arXiv:1405.0599v2 [math.CO] 9 Sep 2014

Multipodal Structure and Phase Transitions in Large Constrained Graphs

Richard Kenyon* Charles Radin† Kui Ren ‡ Lorenzo Sadun§

September 10, 2014

Abstract

We study the asymptotics of large, simple, labeled graphs constrained by the densities of edges and of k-star subgraphs, $k \geq 2$ fixed. We prove that under such constraints graphs are "multipodal": asymptotically in the number of vertices there is a partition of the vertices into $M < \infty$ subsets $V_1, V_2, \ldots, V_M$, and a set of well-defined probabilities g_{ij} of an edge between any $v_i \in V_i$ and $v_j \in V_j$. For $2 \leq k \leq 30$ we determine the phase space: the combinations of edge and k-star densities achievable asymptotically. We also derive the phase space for the triple of densities: edge, 2-star and 3-star. For these models there are special points on the boundary of the phase space with nonunique asymptotic (graphon) structure; for the 2-star model we prove that the nonuniqueness extends to entropy maximizers in the interior of the phase space.

1 Introduction

We study the asymptotics of large, simple, labeled graphs constrained to have certain fixed *densities* t_j of subgraphs H_j, $1 \leq j \leq \ell$ (see definition below). To study the asymptotics we use the graphon formalism of Lovász *et al* [LS1, LS2, BCLSV, BCL, LS3] and the large deviations theorem of Chatterjee and Varadhan [CV], from which one can reduce the analysis to the study of the graphons which maximize the entropy subject to the density constraints, as in the previous works [RS1, RS2, RRS].

Most of this work considers the simple cases, called *k-star models*, in which $\ell = 2$, H_1 is an edge, and H_2 is a "k-star": $k \geq 2$ edges with a common vertex. For these models we prove that all graphons which maximize the entropy, subject to any realizable values of the density constraints, are "multipodal": there is a partition of the vertices into $M < \infty$

*Department of Mathematics, Brown University, Providence, RI 02912; rkenyon@math.brown.edu
†Department of Mathematics, University of Texas, Austin, TX 78712; radin@math.utexas.edu
‡Department of Mathematics, University of Texas, Austin, TX 78712; ren@math.utexas.edu
§Department of Mathematics, University of Texas, Austin, TX 78712; sadun@math.utexas.edu

subsets $V_1, V_2, \ldots, V_M$, and a set of well-defined probabilities g_{ij} of an edge between any $v_i \in V_i$ and $v_j \in V_j$. In particular the optimizing graphons are piecewise constant, attaining only finitely many values.

For any finite set of constraining subgraphs H_j, $1 \leq j \leq \ell$, one can consider the *phase space*, the subset of the unit cube in $\mathbb{R}^\ell$ consisting of accumulation points of all densities $t = (t_1, \ldots t_\ell)$ achievable by finite graphs. The phase space for the 2-star model, and all graphons corresponding to (that is, with densities on) boundary points of the phase space, were derived in [AK]; we derive this for the 3-star model and the *taco model* for which $\ell = 3$ and there are constraints on three subgraphs: edges, 2-stars (sometimes called cherries) and 3-stars (sometimes called claws). We also give evidence of the phase space and bounding graphons for k-star models with $2 \leq k \leq 30$. In these models there are distinguished points on the boundary for which the graphon is not unique. For the 2-star model we prove that this nonuniqueness extends to the entropy maximizer in the interior of the phase space.

Extremal graph theory, the study of the boundaries of the phase spaces of networks, has a long and distinguished history; see for instance [B]. Few examples have been solved, the main ones corresponding to two constraints: edges and the complete graph K_p on $p \geq 3$ vertices [R], which includes the edge/triangle model discussed below. In these examples the optimal graphs are multipodal. However more recently [LS3] examples of 'finitely forced graphons' were found: these are (typically non-multipodal) graphons uniquely determined by the values of (finitely many) subgraph densities. Although we are mainly interested in entropy-maximizing graphons for constraints in the *interior* of the phase space, where one can define phases, (see [RS1] and references therein), these results for graphons corresponding to boundary points are clearly relevant to our study, and will be discussed in Section 8 and the Conclusion.

The significance of multipodal entropy optimizers emerged in a series of three papers [RS1, RS2, RRS] on a model with constraints on edges and triangles, rather than edges and stars. In the edge/triangle model evidence, but not proof, was given that entropy optimizers were M-podal throughout the whole of the phase space, M growing without bound as the two densities approach 1. Here we *prove* that all optimizers are M-podal, with a uniform bound on M, in any k-star model.

A related but different family of models consists of exponential random graph models (ERGMs): see for instance [N, Lov] and the many references therein. In physics terminology the models in [RS1, RS2, RRS] and this paper are microcanonical whereas the ERGMs based on the same subgraph densities are the corresponding grand canonical versions or ensembles. In distinction with statistical mechanics with short range forces [Ru, TET], here the microcanonical and grand canonical ensembles are inequivalent [RS1]. This has important implications on the notion of *phases* in random graph models. All information about the grand canonical ensemble can be derived from the microcanonical ensemble, but not vice-versa. In the Conclusion below we discuss the extent of the loss of information in ERGMs as compared with microcanonical models. Continuing the analogy with statistical mechanics we also discuss the relevance of multipodal structure in the study of emergent phases in all such parametric families of large graphs, as vertex number grows.

2 Notation and background

Fix distinct positive integers $k_1, \ldots, k_\ell$, $\ell \geq 2$, and consider simple (undirected, with no multiple edges or loops) graphs G with vertex set $V(G)$ of labeled vertices, and for each $k = k_i$, the k-star set $T_k(G)$, the set of graph homomorphisms from a k-star into G. We assume $k_1 = 1$ so the k_1-star is an edge. Let $n = |V(G)|$. The *density* of a subgraph H refers to the relative fraction of maps from $V(H)$ into $V(G)$ which preserve edges: the k-star density is

$$t_k(G) \equiv \frac{|T_k(G)|}{n^{k+1}}. \tag{1}$$

For $\alpha > 0$ and $\tau = (\tau_1, \cdots \tau_\ell)$ define $Z_\tau^{n,\alpha}$ to be the number of graphs with densities

$$t_{k_i}(G) \in (\tau_i - \alpha, \tau_i + \alpha),\ 1 \leq i \leq \ell. \tag{2}$$

We sometimes denote τ_1 by ϵ and $T_1(G)$ by $E(G)$.

Define the *(constrained) entropy density* s_τ to be the exponential rate of growth of $Z_\tau^{n,\alpha}$ as a function of n:

$$s_\tau = \lim_{\alpha \downarrow 0} \lim_{n \to \infty} \frac{\ln(Z_\tau^{n,\alpha})}{n^2}. \tag{3}$$

The double limit defining the entropy density s_τ is known to exist [RS1]. To analyze it we make use of a variational characterization of s_τ, and for this we need further notation to analyze limits of graphs as $n \to \infty$. (This work was recently developed in [LS1, LS2, BCLSV, BCL, LS3]; see also the recent book [Lov].) The (symmetric) adjacency matrices of graphs on n vertices are replaced, in this formalism, by symmetric, measurable functions $g : [0,1]^2 \to [0,1]$; the former are recovered by using a partition of $[0,1]$ into n consecutive subintervals. The functions g are called graphons.

For a graphon g define the *degree function* $d(x)$ to be $d(x) = \int_0^1 g(x,y)dy$. The k-star density of g, $t_k(g)$, then takes the simple form

$$t_k(g) = \int_0^1 d(x)^k \, dx. \tag{4}$$

Finally, the *Shannon entropy density* of g is

$$\mathcal{S}(g) = \frac{1}{2} \int_{[0,1]^2} S[g(x,y)] \, dxdy, \tag{5}$$

where S is the Shannon entropy function

$$S(w) = -w \log w - (1-w) \log(1-w). \tag{6}$$

The following is a minor variant of a result in [RS1] (itself an adaption of a proof in [CV]):

Theorem 2.1 (The Variational Principle.)**.** *For any feasible set τ of values of the densities $t(g)$ we have $s_\tau = \max[\mathcal{S}(g)]$, where the maximum is over all graphons g with $t(g) = \tau$.*

(Some authors use instead the *rate function* $I(g) \equiv -\mathcal{S}(g)$, and then minimize I.) The existence of a maximizing graphon $g = g_\tau$ for any constraint $t(g) = \tau$ was proven in [RS1], again adapting a proof in [CV]. If the densities are that of one or more k-star subgraphs we refer to this maximization problem as a *star model*, though we emphasize that the result applies much more generally [RS1].

We consider two graphs *equivalent* if they are obtained from one another by relabeling the vertices. For graphons, the analogous operation is applying a measure-preserving map ψ of $[0,1]$ into itself, replacing $g(x,y)$ with $g(\psi(x),\psi(y))$, see [Lov]. The equivalence classes of graphons under relabeling are called *reduced graphons*, and on this space there is a natural metric, the *cut metric*, with respect to which graphons are equivalent if and only if they have the same subgraph densities for all possible finite subgraphs [Lov]. In the remaining sections of the paper, to simplify the presentation we will sometimes use a convention under which a graphon may be said to have some property if there is a relabeling of the vertices such that it has the property.

The graphons which maximize the constrained entropy tell us what 'most' or 'typical' large constrained graphs are like: if g_τ is the only reduced graphon maximizing $\mathcal{S}(g)$ with $t(g) = \tau$, then as the number n of vertices diverges and $\alpha_n \to 0$, exponentially most graphs with densities $t_i(G) \in (\tau_i - \alpha_n, \tau_i + \alpha_n)$ will have reduced graphon close to g_τ [RS1]. This is based on large deviations from [CV].

3 Multipodal Structure

Usually our graphons g have a constraint on edge density, which is the integral of g over $[0,1]^2$, so we treat g as a way of assigning this conserved quantity, which for intuitive purposes we term 'mass', to the various regions of $[0,1]^2$.

Except where otherwise indicated, in this section we restrict attention to a k-star model for fixed $k \geq 2$. Note that the values of $\tau(g)$ are determined by the degree function $d(x)$, as $t_k = \int_0^1 d(x)^k dx$. We first prove a general result for graphons with constraints only on their degree function. (See Theorem 1.1 in [CDS] for a stronger result but with stronger hypotheses.)

Theorem 3.1. *If the degree function $d(x)$ is monotonic nondecreasing and if the graphon $g(x,y)$ maximizes the entropy among graphons with the same degree function, then $g(x,y)$ is doubly monotonic nondecreasing, that is, outside a set of measure zero, g is monotonic nondecreasing in each variable.*

Before providing a rigorous proof, consider the following heuristic. Suppose that $d(x)$ is monotonically nondecreasing and that $g(x_1,y_1) > g(x_2,y_1)$ for some $x_1 < x_2$. Since $d(x_1) \leq d(x_2)$, there must be some other value of y, say y_2, such that $g(x_1,y_2) < g(x_2,y_2)$. But then moving mass from (x_1,y_1) to (x_2,y_1) and moving the same amount of mass from (x_2,y_2) to (x_1,y_2) (and likewise moving mass from (y_1,x_1) to (y_1,x_2) and from (y_2,x_2) to (y_2,x_1) to preserve symmetry) will increase the entropy $\mathcal{S}$ while leaving the degree function

$d(x)$ fixed. The problem with this heuristic is that mass is distributed continuously, so we cannot speak of mass "at a point". Instead, we must smear out the changes over sets of positive measure. The following proof is essentially the above argument, averaged over all possible data (x_1, x_2, y_1, y_2).

Proof. For $0 < a, b < 1$, define

$$\begin{aligned} \eta(a,b) &= \int_0^1 \max(g(a,y) - g(b,y), 0)\, dy \\ \kappa(a,b) &= \int_0^1 \max(g(b,y) - g(a,y), 0)\, dy. \end{aligned} \tag{7}$$

Note that $\kappa(a,b) - \eta(a,b) = d(b) - d(a) \geq 0$, and that $g(x,y)$ being doubly monotonic almost everywhere is equivalent to $\eta(a,b)$ being almost everywhere zero when $b \geq a$. For $b \geq a$ let

$$\gamma(y,a,b) = \begin{cases} \eta(a,b), & g(b,y) \geq g(a,y) \\ \kappa(a,b), & g(b,y) < g(a,y) \end{cases}, \tag{8}$$

and for $a > b$ let $\gamma(y,a,b) = \gamma(y,b,a)$. If $g(x,y)$ is doubly monotonic, then $g(b,y) \geq g(a,y)$ for $b \geq a$, so γ is equal to η almost everywhere, which is equal to zero almost everywhere.

Now evolve the graphon g according to the following differential equation:

$$\frac{d}{dt} g_t(x,y) = \int_0^1 \gamma_t(y,x,b)[g_t(b,y) - g_t(x,y)]\, db + \int_0^1 \gamma_t(x,y,b)[g_t(x,b) - g_t(x,y)]\, db. \tag{9}$$

If $g(x,y)$ is doubly monotonic almost everywhere, then this simplifies to $dg_t/dt = 0$. Otherwise, we will show that the flow preserves the degree function $d(x)$ and increases the entropy $\mathcal{S}(g)$, contradicting the assumption that g is an entropy maximizer.

To see that the degree function is unchanged we compute

$$\begin{aligned} \frac{d}{dt} d_t(x) &= \int_0^1 \frac{d}{dt} g_t(x,y)\, dy \\ &= \int_0^1 \int_0^1 \gamma_t(y,x,b)[g_t(b,y) - g_t(x,y)]\, dy\, db \\ &\quad + \int_0^1 \int_0^1 \gamma_t(x,y,b)[g_t(x,b) - g_t(x,y)]\, dy\, db \end{aligned} \tag{10}$$

The second integrand is anti-symmetric in b and y, and so integrates to zero. For the first integral, fix a value of b and write

$$g_t(b,y) - g_t(x,y) = \max(g_t(b,y) - g_t(x,y), 0) - \max(g_t(x,y) - g_t(b,y), 0). \tag{11}$$

Suppose for the moment that $b > x$. Integrating over those values of y for which $g(b,y) \geq g(x,y)$ yields $\eta_t(x,b)\kappa_t(x,b)$, since $\gamma_t(y,x,b) = \eta_t(x,b)$ and the integral of $\max(g_t(b,y) - g_t(x,y), 0)$ gives $\kappa_t(x,b)$. Similarly, integrating over those values of y for which $g_t(b,y) < g_t(x,y)$ gives $-\kappa_t(x,y)\eta_t(x,b)$, so the sum is zero. A similar argument applies when $b < x$,

only with the roles of η_t and κ_t reversed. Since the integral over y is zero for almost every b, the double integral is zero.

To see that $\mathcal{S}(g_t)$ is increasing, we compute

$$\begin{aligned}
\frac{d}{dt}\mathcal{S}(g_t) &= \iint S'(g_t(x,y))\frac{d}{dt}g_t(x,y)\,dx\,dy \\
&= \iint S'(g_t(x,y)\,dx\,dy \int_0^1 \gamma_t(y,x,b)[g_t(b,y)-g_t(x,y)] \\
&\qquad\qquad +\gamma_t(x,y,b)[g_t(x,b)-g_t(x,y)]\,db \\
&= 2\iint S'(g_t(x,y))\,dx\,dy \int \gamma_t(x,y,b)[g_t(x,b)-g_t(x,y)]\,db \\
&= \iiint \gamma_t(x,y,b)[g_t(x,b)-g_t(x,y)][S'(g_t(x,y))-S'(g_t(x,b))]\,dx\,dy\,db \quad (12)
\end{aligned}$$

However, $[g_t(x,b)-g_t(x,y)][S'(g_t(x,y))-S'(g_t(x,b))]$ is strictly positive when $g_t(x,b)\neq g_t(x,y)$, thanks to the concavity of the function S. So $d\mathcal{S}(g_t)/dt$ is non-negative, and is strictly positive unless $\gamma_t(x,y,b)$ is zero for (almost) all triples (x,y,b) for which $g_t(x,b)\neq g_t(x,y)$. Since the vanishing of γ_t is equivalent to the double monotonicity of g_t, a maximizing g must be doubly monotonic. □

If $g(x,y)$ is doubly monotonic almost everywhere, we can adjust it on a set of measure zero to be doubly monotonic everywhere. Just take the adjusted value $\tilde{g}(a,b)$ to be the essential supremum of $g(x,y)$ over all (x,y) with $x<a$ and $y<b$. Since changes over sets of measure zero have no effect on the integrals of g, we can therefore assume hereafter that $g(x,y)$ is doubly monotonic whenever g is an entropy maximizer.

As a first corollary to this theorem, we have the following result generalizing slightly a result of [CDS]:

Proposition 3.2. *If g maximizes the entropy among graphons with the same degree function, and if the degree function d takes only M values, then g is M-podal.*

Proof. Rearrange vertices to make $d(x)$ to be monotone increasing. If d is constant on some interval $[x_1,x_2]$ then for any y, $g(x,y)$ is also constant for $x\in[x_1,x_2]$: if $g(x_2,y)>g(x_1,y)$ then there would have to be some y' such that $g(x_2,y')<g(x_1,y')$ to assure that $d(x_1)=d(x_2)$, contradicting monotonicity of g. □

As a second corollary, we obtain a strong continuity result:

Proposition 3.3. *If g maximizes the entropy among graphons with a given degree function, then g is continuous almost everywhere.*

Proof. We may assume that g is doubly monotonic. This implies that g is monotonic (and of course bounded) on any line $y=x+c$ of slope 1. dg is then a bounded measure on this line, whose discrete part is supported on a countable number of points. In other words, $g(x,x+c)$ can only have a countable number of jump discontinuities as a function of x, and

is otherwise continuous. By Fubini's theorem, $g(x,y)$ must then be continuous in the $(1,1)$ direction for almost all (x,y). But if g is continuous in the $(1,1)$ direction at a point (a,b), then for each $\epsilon > 0$ we can find a $\delta > 0$ such that $g(a+\delta, b+\delta)$ and $g(a-\delta, b-\delta)$ are both within ϵ of $g(a,b)$. But since g is doubly monotonic, $g(a-\delta, b-\delta) \le g(x,y) \le g(a+\delta, b+\delta)$ for all $x \in (a-\delta, a+\delta)$ and all $y \in (b-\delta, b+\delta)$, so g is continuous at (a,b). □

The upshot of this proposition is that the value of g at a generic point (a,b), and the functional derivatives $\delta\mathcal{S}/\delta g$ and $\delta t_k/\delta g$ at (a,b), control the values of these functions in a neighborhood of (a,b). (By the functional derivative $\delta t_k/\delta g$ we mean the function such that $\int \delta t_k/\delta g(x,y)\,dxdy$ is the coefficient of the linear term in the expansion of $t_k(g+\delta g)$.) We can therefore do functional calculus computations at points (a,b) and (c,d), and then speak of moving mass from a neighborhood of (a,b) to a neighborhood of (c,d). In other words, once we have (almost everywhere) continuity, informal arguments such as those preceding the proof of Theorem 3.1 can be used directly.

Here is our main theorem.

Theorem 3.4. *For the k-star model, any graphon g which maximizes the entropy $\mathcal{S}(g)$ and is constrained by $t(g) = \tau$, is M-podal for some $M < \infty$.*

Proof. When τ is on the boundary of the phase space Theorem 4.1, below, indicates that the graphon is either 1-podal (on the lower boundary) or ≤ 3 podal on the top boundary. So for the remainder of the proof we assume τ is in the interior.

Lemma 3.5. *For the k-star model, let g be a graphon that maximizes the entropy $\mathcal{S}(g)$ subject to the constraints $t(g) = \tau$, where τ lies in the interior of the phase space of possible densities. Then there exist constants β_1, β_2 such that the Euler-Lagrange equation*

$$\beta_1 + \beta_2 d(x)^{k-1} + \beta_2 d(y)^{k-1} = \ln\left[\frac{1}{g(x,y)} - 1\right] \tag{13}$$

holds for almost every (x,y). Furthermore, the constants β_1, β_2 are uniquely defined.

Proof of lemma. First note that g cannot take values in $\{0,1\}$ only, since such a graphon would have zero entropy, and for each τ in the interior of the phase space it is easy to construct a graphon (even bipodal, see section 6) with positive entropy.

We claim that $g \in (0,1)$ on a set of full measure. Suppose otherwise. Then (by double monotonicity) $g(x,y)$ is either 1 on a neighborhood of $(1,1)$ or is zero on a neighborhood of $(0,0)$ (or both). My moving ϵ mass from a neighborhood of $(1,1)$ to a neighborhood of $(0,0)$, we can increase the entropy by order $\epsilon \ln(1/\epsilon)$, while leaving the edge density fixed. In the process, we will decrease t_k, since the monotonicity of $d(x)$ implies that the functional derivative $\delta t_k/\delta g = (k/2)(d(x)^{k-1} + d(y)^{k-1})$ is greater near $(1,1)$ than near $(0,0)$. However, we claim that we can restore the value of t_k by moving mass within the region R_0 where $g \in (0,1)$, at a cost in entropy of order ϵ. Since $\epsilon \ln(1/\epsilon) \gg \epsilon$, for sufficiently small ϵ the combined move will increase entropy while leaving t_1 and t_k fixed, which is a contradiction.

The details of the second movement of mass depend on whether $\delta t_k/\delta g$ is constant on R_0 or not. If $\delta t_k/\delta g$ is not constant, we can restrict attention to a slightly smaller region $\tilde{R}_0$ where g is bounded away from 0 and 1, and then move mass from a portion of $\tilde{R}_0$ where $\delta t_k/\delta g$ is smaller to a portion where $\delta t_k/\delta g$ is larger. This will increase t_k to first order in the amount of mass moved. Since g is bounded away from 0 and 1 on $\tilde{R}_0$, the change in the entropy is bounded by a constant times the amount of mass moved, as required.

If $\delta t_k/\delta g$ is constant on R_0, then it is constant on a rectangle within R_0. But the only way for $d(x)^{k-1} + d(y)^{k-1}$ to be constant on a rectangle is for $d(x)$ and $d(y)$ to be constant for (x, y) in that rectangle. By Theorem 3.1, this implies that $g(x, y)$ is constant on the rectangle. Moving mass within the rectangle will then change neither the entropy nor t_k to first order, but will change both (with t_k increasing and the entropy decreasing) to second order. So by moving an amount of mass of order $\sqrt{\epsilon}$, we can restore the value of t_k at an $O(\epsilon)$ cost in entropy. This proves our assertion that $g(x, y) \in (0, 1)$ on a set of full measure.

Next we note that the degree function $d(x)$ must take on at least two values, since otherwise we would be on the lower boundary of the phase space, with $t_k = t_1^k$. This means that the functional derivative

$$\frac{\delta t_k}{\delta g}(x, y) = \frac{k}{2}\left(d(x)^{k-1} + d(y)^{k-1}\right) \tag{14}$$

is not a constant function. If $\delta\mathcal{S}/\delta g(x, y)$ cannot be written as a linear combination of $\delta t_1/\delta g(x, y) = 1$ and $\delta t_k/\delta g$, then we can find three points $p_i = (x_i, y_i)$ such that g is continuous at each point, and such that the matrix

$$\begin{pmatrix} 1 & 1 & 1 \\ \frac{\delta t_k}{\delta g}(p_1) & \frac{\delta t_k}{\delta g}(p_2) & \frac{\delta t_k}{\delta g}(p_3) \\ \frac{\delta \mathcal{S}}{\delta g}(p_1) & \frac{\delta \mathcal{S}}{\delta g}(p_2) & \frac{\delta \mathcal{S}}{\delta g}(p_3) \end{pmatrix} \tag{15}$$

is invertible. But then, by adjusting the amount of mass near each p_i (and near the reflected points $p_i' = (y_i, x_i)$), we can independently vary t_1, t_k and $\mathcal{S}$ to first order. By the inverse function theorem, we can then increase $\mathcal{S}$ while leaving t_1 and t_k fixed, which is a contradiction. Thus $\delta\mathcal{S}/\delta g$ must be a linear combination of $\delta t_1/\delta g$ and $\delta t_k/\delta g$, which is equation (13). Since $\delta t_1/\delta g$ and $\delta t_k/\delta g$ are linearly independent, the coefficients are unique. □

Continuing the proof of Theorem 3.4, solving (13) for $g(x, y)$ gives

$$g(x, y) = \frac{1}{1 + \exp[\beta_1 + \beta_2 d(x)^{k-1} + \beta_2 d(y)^{k-1}]}, \tag{16}$$

and integrating with respect to y gives

$$d(x) = \int_0^1 \frac{dy}{1 + \exp[\beta_1 + \beta_2 d(x)^{k-1} + \beta_2 d(y)^{k-1}]}. \tag{17}$$

Let $d(x)$ be any solution of (17), let z be a real variable, and consider the function

$$F(z) = z - \int_0^1 \frac{dy}{1 + \exp[\beta_1 + \beta_2 z^{k-1} + \beta_2 d(y)^{k-1}]}, \tag{18}$$

where the function $d(y)$ is treated as given. By equation (17), all actual values of $d(x)$ are roots of $F(z)$.

The second term in (18) is an analytic function of z, as follows.

Write $W = \beta_2 z^{k-1}$ and $Y = \beta_2 d(y)^{k-1}$ then the integral is

$$\int \frac{d\mu(Y)}{1 + \exp(\beta_1 + W + Y)}, \tag{19}$$

the convolution of an analytic function of W with an integrable measure $\mu(Y)$. Since the Fourier transform of an analytic function decays exponentially at infinity and the Fourier transform of an integrable measure is bounded, the Fourier transform of the convolution decays exponentially at infinity, so the convolution itself is an analytic function of W. Since W is an analytic function of z, $F(z)$ is an analytic function of z.

Note that $F(z)$ is strictly negative for $z \leq 0$ and strictly positive for $z \geq 1$. Being analytic and not identically zero, $F(z)$ can only have finitely many roots in any compact interval. By Rolle's Theorem any accumulation point of the roots would have to be an accumulation point of the roots of $F'(z)$, $F''(z)$, etc. So all derivatives of F would have to vanish at that point, making the Taylor series around it identically zero. In particular $F(z)$ can only have finitely many roots between 0 and 1, implying there are only finitely many values of $d(x)$. By Proposition 3.2, the graphon g is M-podal. Note that the roots of $F(z)$ are not necessarily values of $d(x)$, so this construction only gives an upper bound to the actual value of M. □

Let c_j be the measure of the set $\{x \in [0,1] \mid d(x) = d_j\}$. We can apply a measure-preserving transformation so that $d(x) = d_1$ on $[0, c_1]$, $d(x) = d_2$ on $[c_1, c_1 + c_2]$, etc. If x is in the j-th interval and y is in the m-th interval, then

$$g(x,y) = \frac{1}{1 + \exp[\beta_1 + \beta_2(d_j^{k-1} + d_m^{k-1})]}. \tag{20}$$

Thus g is piecewise constant. A graph with N vertices corresponding to this graphon will have approximately Nc_j vertices with degree close to Nd_j for each $j \in \{1, 2, \ldots, M\}$. There will be approximately $N^2 c_j c_m / (1 + \exp[\beta_1 + \beta_2(d_j^{k-1} + d_m^{k-1})])$ edges between vertices in cluster j and vertices in cluster m (if $j \neq m$, half that if $j = m$), and these edges will be statistically independent of one another.

The previous proof actually showed more than that optimal graphons are multipodal. We showed that the possible values of $d(x)$ are roots of the function $F(z)$ defined in (18). This allows us to prove the following refinement of Theorem 3.4

Theorem 3.6. *For any k-star model, there exists a fixed M such that all entropy-maximizing graphons are m-podal with $m \leq M$.*

Proof. If $f(z)$ is an analytic function on a compact interval with m roots (counted with multiplicity), then any C^m-small perturbation of $f(z)$ will also have at most m roots on the interval. Thus if g is a graphon with associated function $F(z)$ with m roots, and if $\tilde{g}$ is an

optimal graphon whose degree function is L^1-close to that of g, then the associated function $\tilde{F}(z)$ of $\tilde{g}$ will have at most m roots, and $\tilde{g}$ will be at most m-podal.

Suppose there is no universal bound M on the podality of optimal graphons on R. Let $g_1, g_2, \ldots$ be a sequence of optimal graphons (perhaps with different values of τ) with the podality going to infinity, and let $F_i(z)$ be the associated functions for these graphons. Since the space of graphons is compact, there is a subsequence that converges to a graphon g_∞. The associated function $F_\infty(z)$ of g_∞ is analytic, and so has only a finite number M_∞ of roots. But then for large i, $F_i(z)$ has at most M_∞ roots and g_i is at most M_∞-podal, which is a contradiction. □

Corollary 3.7. *The entropy function s_τ is piecewise analytic as a function of τ.*

Proof. Since all optimal graphons are M-podal (where we can treat m-podal structures with $m < M$ as degenerate M-podal structures), the entropy and k-star densities of M-podal graphons are analytic functions of the parameters c_i and g_{ij}. We then have an $M + \binom{M+1}{2}$ dimensional variational problem involving analytic functions. The Euler-Lagrange equations for stationary points of the entropy are then finite systems of analytic equations, whose solutions must be analytic functions of the parameter τ. Consequently, wherever the maximizing graphon varies continuously with τ, it must vary analytically, as must the entropy. The only places where s_τ is not real-analytic is along codimension-1 "phase transition" surfaces where the maximization problem has two (or more) inequivalent solutions.

□

We end this section with an argument which displays the use of the notion of phase. By definition a phase is a connected open subset of the phase space in which the entropy s_τ is analytic. (The connection with statistical mechanics is discussed in the Conclusion.) The following only simplifies one step in the proof of our main result, Theorem 3.4, but it shows how the notion of phase can be relevant.

Theorem 3.8. *For any microcanonical model let g be a graphon which maximizes the Shannon entropy $\mathcal{S}(g)$ subject to the constraints $t(g) = \tau$, where τ lies in the interior of the phase space of possible densities and s_τ is differentiable at τ. Then the set $A = g_0^{-1}(\{0, 1\})$ has measure zero.*

Proof of theorem. Define g_ϵ by moving the value of g_0 on A by ϵ, $0 < \epsilon < 1$.

From their definitions the densities satisfy $t(g_\epsilon) = t(g_0) + O(\epsilon)$. By integrating over A we see that the Shannon entropy satisfies

$$\mathcal{S}(g_\epsilon) = \mathcal{S}(g_0) - m(A)\epsilon \ln(\epsilon), \tag{21}$$

so noting that $s_{t(g_\epsilon)} \geq \mathcal{S}(g_\epsilon)$ and $s_{\tau_0} = \mathcal{S}(g_0)$ we get

$$s_{t(g_\epsilon)} \geq s_{\tau_0} - m(A)\epsilon \ln(\epsilon). \tag{22}$$

From differentiability, as the vector $\alpha \to 0$

$$s_{\tau_0+\alpha} = s_{\tau_0} + O(||\alpha||). \tag{23}$$

So as $\epsilon \to 0$ we have a contradiction with (23) unless $m(A) = 0$, which concludes the proof.

□

4 Phase space

We now consider k-star models with $2 \leq k \leq 30$. The phase space is the set of those $(\epsilon, \tau) \subset [0,1]^2$ which are accumulation points of the values of pairs (edge density, k-star density) for finite graphs. The lower boundary (minimum of τ given ϵ) is easily seen to be the Erdős-Rényi curve: $\tau = \epsilon^k$, since Hölder's inequality gives

$$\tau^{1/k} = \|d(x)\|_k \geq \|d(x)\|_1 = \epsilon. \tag{24}$$

We now determine the upper curve. This was determined for $k = 2$ in [AK], and perhaps was published for higher k though we do not know a reference.

We are looking for the graphon which maximizes $t(g)$ for fixed $e(g) = \epsilon$, and this time arrange the points of the line so that $d(x)$ is monotonically decreasing. As in the proof of Theorem 3.4 we can assume that g is monotonic (this time, decreasing) in both coordinates.

We call a graphon a *g-clique* if it is bipodal of the form

$$g(x,y) = \begin{cases} 1 & x < c \text{ and } y < c \\ 0 & \text{otherwise} \end{cases} \tag{25}$$

and a *g-anticlique* if it is of the form

$$g(x,y) = \begin{cases} 0 & x > c \text{ and } y > c \\ 1 & \text{otherwise.} \end{cases} \tag{26}$$

Theorem 4.1. *For fixed $e(g) = \epsilon$, and any $2 \leq k \leq 30$, any graphon that maximizes the k-star density is equivalent to a g-clique or g-anticlique.*

G-cliques always have $c = \sqrt{\epsilon}$ and k-star density $\epsilon^{(k+1)/2}$. G-anticliques have $c = 1 - \sqrt{1-\epsilon}$ and k-star density

$$c + c^k - c^{k+1} = [1 - \sqrt{1-\epsilon}][1 + \sqrt{1-\epsilon}(1 - \sqrt{1-\epsilon})^{k-1}]. \tag{27}$$

For ϵ small, the g-anticlique has k-star density $\epsilon/2 + O(\epsilon^2)$, which is greater than $\epsilon^{(k+1)/2}$. For ϵ close to 1, however, the g-clique has a higher k-star density than the g-anticlique.

While our proof only covers k up to 30, we conjecture that the result holds for all values of k. The only difficulty in extending to all k is comparing the t_k-value for a clique, anticlique and the "tripodal anticlique" in the last figure of Figure 7.

Corollary 4.2. *For $2 \le k \le 30$, the upper boundary of the phase space is*

$$\tau = \begin{cases} [1-\sqrt{1-\epsilon}][1+\sqrt{1-\epsilon}(1-\sqrt{1-\epsilon})^{k-1}] & \epsilon \le \epsilon_0 \\ \epsilon^{(k+1)/2} & \epsilon \ge \epsilon_0 \end{cases} \tag{28}$$

where $\epsilon = \epsilon_0$ where the two branches of $\tau = \tau(\epsilon)$ cross.

For $k = 2$ the crossing point is $\epsilon_0 = 1/2$, for $k = 3$ it is $\epsilon_0 = 3/4$, and as $k \to \infty$ it approaches 1. The boundary of the phase space for the 2-star model is shown in Fig. 1.

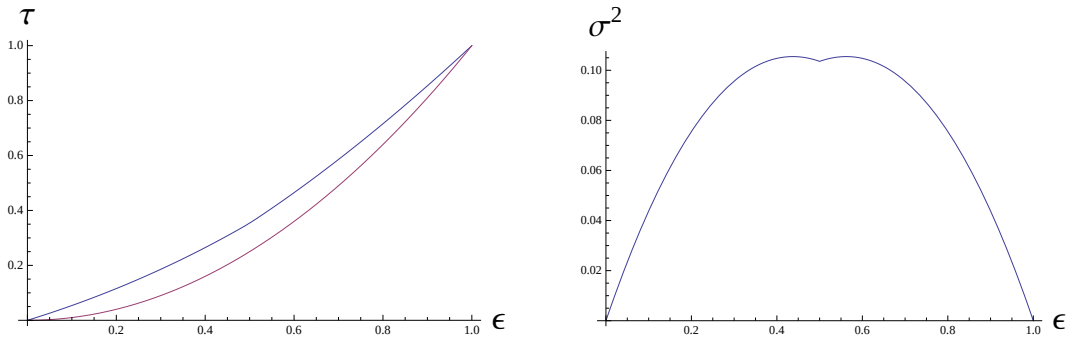

Figure 1: Boundary of the phase space for the 2-star case. Left: true phase boundary; Right: Plot of ϵ versus $\sigma^2 = \tau - \epsilon^2$; in this case the lower boundary becomes the x-axis.

4.1 Proof of Theorem 4.1

The proof has four steps:

1. Showing that $g(x, y)$ only takes on the values 0 and 1.
2. Showing that the boundary between the $g = 1$ and $g = 0$ regions is a zig-zag sequence of vertical and horizontal segments.
3. Showing that there cannot be more than three interior corners to the zig-zag.
4. Showing that there is only one interior corner. Depending on whether the zig-zag hits the boundary of the square at the coordinate axes or at the lines $x = 1$ and $y = 1$, the result is then either a g-clique or a g-anticlique. (This is the only step that uses $k \le 30$.)

4.1.1 Step 1: Showing we have a 0-1 graphon

For illustrative clarity, we show the first step for $k = 2$ and then generalize.

The variational equation for maximizing $t_k(g)$ while fixing $t_1(g)$ is (see (13) without the '$\mathcal{S}$' term)

$$d(x) + d(y) = \lambda \tag{29}$$

for some unknown constant λ. This should not be taken too literally, since it only applies where $0 < g(x,y) < 1$. When $g(x,y) = 0$ we have $d(x) + d(y) \leq \lambda$, and when $g(x,y) = 1$ we have $d(x) + d(y) \geq \lambda$. Since $d(x) + d(y)$ is a decreasing function of both x and y, this means that there is a strip (of possibly zero thickness) running roughly from the northwest corner of $[0,1]^2$ to the southeast corner, where $g(x,y)$ is strictly between 0 and 1 and where $d(x) + d(y) = \lambda$. All points to the southwest of this strip have $g(x,y) = 1$, and all points to the northeast have $g(x,y) = 0$. (The boundaries of the strip are necessarily monotone paths).

We claim that the strip has zero thickness, and hence is actually a boundary path between the $g = 1$ and the $g = 0$ zone. To see this, suppose that the strip contains a small ball, and hence contains a small rectangle. Since $d(x) + d(y)$ is constant on this rectangle, $d(x)$ is constant and $d(y)$ is constant. We can then increase the 2-star density (i.e. second moment of $d(x)$) by moving some mass from right to left, and from top to bottom on the mirror image region, as in the proof of Theorem 3.4. As in that proof, this move increases $t_k(g)$ to second order in the amount of mass moved. Since we assumed that our graphon was a maximizer (and not just a stationary point), we have a contradiction.

If $k > 2$, then the variational equations are $d(x)^{k-1} + d(y)^{k-1} = \lambda$ rather than $d(x) + d(y) = \lambda$. Otherwise, the arguments are exactly the same, since $d(x)^{k-1} + d(y)^{k-1}$ is still a monotonic function of both x and y.

4.1.2 Step 2: Showing we have a zig-zag boundary

From the remarks after (29) above, we have that $d(x)^{k-1} + d(y)^{k-1}$ is greater (or equal), at all points (x,y) with $g(x,y) = 1$, than it is at all points where $g(x,y) = 0$. For (x,y) on a non-vertical and non-horizontal part of the boundary between the regions, $d(x) = y$ and $d(y) = x$, by Step 1 above, so those parts of the boundary have to lie on a line $x^{k-1} + y^{k-1} =$constant. But such a curve does not maximize t_k! Changing the curve to a horizontal-vertical zig-zag will increase t_k to second order in the amount of mass moved. (For instance, if $k = 2$ and the boundary was the line $x + y = \lambda$ between $x = a$ and $x = b$, then the boundary consisting of segments running from $(a, \lambda - a)$ to $((a+b)/2, \lambda - a)$ to $([a+b]/2, \lambda - b)$ to $(b, \lambda - b)$ (with corresponding changes near (b,a)) has a greater variation in $d(x)$, and hence a larger value of $t_2 = \int_0^1 d^2(x)dx$.)

Since there are no continuous diagonal segments, the boundary must be a vertical-horizontal zig-zag. The northeast-pointing corners of the zig-zag lie strictly above the curve $x^{k-1} + y^{k-1} = \lambda$ and the southwest pointing corners lie below the curve, so the distance from each corner to the curve gives a lower bound for the size of the next (or previous) segment. This eliminates the possibility of a dense accumulation of zig-zags. We either have a finite number of zigs and zags, or at worst a finite number of accumulation points.

4.1.3 Step 3: Limiting the number of segments

If the number of zig-zag segments is finite, there must be points $0 = c_0 < c_1 < c_2 < \cdots < c_n \leq 1$ such that the zig-zag bounds the graph of the function

$$\begin{cases} c_n & x < c_1 \\ c_{n-1} & c_1 < x < c_2 \\ \vdots & \vdots \\ c_j & c_{n-j} < x < c_{n+1-j} \end{cases} \tag{30}$$

(If there are infinitely many zig-zag segments, then we have a similar expression, only involving an infinite step function.) There are n degrees of freedom if $c_n < 1$ and $n-1$ if $c_n = 1$. The graphon is $(n+1)$-podal if $c_n < 1$ and n-podal if $c_n = 1$. In this step we will show that $n \leq 3$, and that $c_3 = 1$ if $n = 3$. This shows that the optimizing graphon is at most 3-podal.

We compute

$$t_k = \sum_{i=1}^{n} (c_i - c_{i-1}) c_{n+1-i}^k.$$

Suppose that $n > 3$ or infinite, or that $n = 3$ and $c_3 < 1$. Then in the notation of equation (30), there is a subpath consisting of three consecutive vertices $(c_{i-1}, c_j)(c_i, c_j)(c_i, c_{j-1})$ where $c_j < 1$ and all three vertices lie strictly above the diagonal $x = y$. We will show that this leads to a contradiction.

The contribution to t_k involving c_i and c_j is

$$t_k = c_j^k(c_i - c_{i-1}) + c_{j-1}^k(c_{i+1} - c_i) + c_i^k(c_j - c_{j-1}) + c_{i-1}^k(c_{j+1} - c_j) + (\text{terms not involving } c_i, c_j).$$

The contribution to $e = t_1$ is

$$t_1 = (c_i - c_{i-1})(c_j - c_{j-1}) + (\text{terms not involving } c_i, c_j).$$

We can rewrite

$$t_k = (c_i - c_{i-1})(c_j^k - c_{j-1}^k) + (c_j - c_{j-1})(c_i^k - c_{i-1}^k) + (\text{terms not involving } c_i, c_j).$$

Let $w = \frac{c_j}{c_{j-1}} \in (1, \frac{c_{j+1}}{c_{j-1}})$ and, if $c_{i-1} \neq 0$, write $z = \frac{c_i}{c_{i-1}} \in (1, \frac{c_{i+1}}{c_{i-1}})$ (we deal with the case $c_{i-1} = 0$ below). Then

$$t_1 = c_{i-1}c_{j-1}(z-1)(w-1) + (\text{terms not involving } c_i, c_j).$$

Grouping terms not involving z, w this becomes

$$(z-1)(w-1) = C$$

for some positive constant C. We also have

$$t_k = c_{i-1}c_{j-1}^k(z-1)(w-1)\frac{w^k - 1}{w-1} + c_{j-1}c_{i-1}^k(z-1)(w-1)\frac{z^k - 1}{z-1} + (\text{terms not involving } c_i, c_j),$$

which we can rewrite (using $(z-1)(w-1)=C$)

$$t_k = A\frac{z^k-1}{z-1} + B\frac{w^k-1}{w-1} + C_2.$$

However for $k \geq 2$ one easily sees that for $A, B, C > 0$ the function

$$A\frac{z^k-1}{z-1} + B\frac{w^k-1}{w-1}$$

is strictly convex on the curve $(z-1)(w-1)=C$ (for example by substituting $x = z-1, y = w-1$), so the maximum value of t_k is attained at one of the endpoints $z=1,\ w = c_{j+1}/c_{j-1}$ or $z = c_{i+1}/c_{i-1},\ w = 1$. So either c_i merges with c_{i+1} or c_j merges with c_{j+1}, so our presumed graphon with $n > 3$ (or with $n = 3$ and $c_3 < 1$) did not actually maximize t_k.

4.1.4 Step 4: Eliminating $n = 3, c_3 = 1$ and $n = 2, c_3 < 1$.

For $n = 3$, $c_3 = 1$, the edge and k-star densities are:

$$t_1 = c_1 + (c_2 - c_1)c_2 + (1-c_2)c_1 = 2c_1 - 2c_1c_2 + c_2^2 \qquad t_k = c_1 + (c_2-c_1)c_2^k + (1-c_2)c_1^k. \quad (31)$$

Taking derivatives with respect to c_j gives:

$$\begin{array}{ll} \partial_1 t_1 = 2(1-c_2); & \partial_2 t_1 = 2(c_2 - c_1); \\ \partial_1 t_k = 1 - c_2^k + k(1-c_2)c_1^{k-1}; & \partial_2 t_k = (c_2^k - c_1^k) + k(c_2-c_1)c_2^{k-1}. \end{array} \quad (32)$$

An important quantity is the ratio $r_j = 2\frac{\partial_j t_k}{\partial_j t_1}$. For $j = 1, 2$, this works out to

$$\begin{aligned} r_1 &= kc_1^{k-1} + \frac{c_3^k - c_2^k}{c_3 - c_2} \\ r_2 &= kc_2^{k-1} + \frac{c_2^k - c_1^k}{c_2 - c_1} \end{aligned} \quad (33)$$

We imagine c_1 and c_2 evolving in time so as to keep t_1 fixed, with $\dot{c}_1 = c_2 - c_1,\ \dot{c}_2 = c_2 - c_3$ and $c_3 = 1$ fixed (note that this satisfies $\dot{t}_1 = 0$ by (32)). We must have $r_1 = r_2$, or else $\dot{t}_k$ would be nonzero. We will show that the $d(r_1 - r_2)/dt > 0$, and hence that t_k increases to second order. We compute:

$$\begin{aligned} \dot{r}_1 - \dot{r}_2 &= k(k-1)[c_1^{k-2}\dot{c}_1 - c_2^{k-2}\dot{c}_2] \\ &\quad + k\left[-\frac{c_1^{k-1}\dot{c}_2}{c_3 - c_2} - \frac{c_2^{k-1}\dot{c}_2}{c_2 - c_1} + \frac{c_1^{k-1}\dot{c}_1}{c_2 - c_1}\right] \\ &\quad + \frac{c_2^k - c_1^k}{c_2 - c_1}\left[\frac{\dot{c}_2}{c_3 - c_2} + \frac{\dot{c}_2}{c_2 - c_1} - \frac{\dot{c}_1}{c_2 - c_1}\right] \end{aligned} \quad (34)$$

Plugging in the values of $\dot{c}_1$ and $\dot{c}_2$ then gives

$$\dot{r}_1 - \dot{r}_2 = k(k-1)\left[c_1^{k-2}(c_2 - c_1) + c_2^{k-2}(1-c_2)\right]$$

$$+k\left[2c_1^{k-1}+c_2^{k-1}\frac{1-c_2}{c_2-c_1}\right]$$
$$-\frac{c_2^k-c_1^k}{c_2-c_1}\left[2+\frac{1-c_2}{c_2-c_1}\right] \tag{35}$$

Let $x=c_1/c_2$ and let $z=1/c_2$. Note that $x<1<z$. Then $\dot{r}_1-\dot{r}_2$ is c_2^{k-1} times

$$\begin{aligned} f(x,z) &:= k(k-1)[x^{k-2}(1-x)+z-1]+k[2x^{k-1}+\frac{z-1}{1-x}]-\left(\frac{1-x^k}{1-x}\right)\left[2+\frac{z-1}{1-x}\right] \\ &= k(k-1)[x^{k-2}(1-x)+z-1]+2[(k-1)x^{k-1}-x^{k-2}-\cdots-1] \\ &\quad +(z-1)[(k-1)+(k-2)x+\cdots+x^{k-2}]. \end{aligned} \tag{36}$$

This is an order $k-1$ polynomial in x and linear in z. Likewise, the constraint $r_1=r_2$ becomes

$$0=F(x,z):=kx^{k-1}+\frac{z^k-1}{z-1}-k-\frac{1-x^k}{1-x}.$$

Note that $F(1,1)=f(1,1)=0$. By implicitly differentiating F we see that $dx/dz=-1$ at $(x,z)=(1,1)$, hence that $df/dz|_{(1,1)}=\frac{3}{2}k(k-1)>0$. Hence $\dot{r}_1-\dot{r}_2$ is positive when x is slightly less than 1. If $f(x,z)$ ever fails to be positive on the set where $F(x,z)=0$ then, by the intermediate value theorem, there is a value of z, and a corresponding value of x, for which $f(x,z)=0$ and $F(x,z)=0$. The intersection of the two degree $k-1$ algebraic curves $F=0$ and $f=0$ would then have to contain a real point (x,z) with $0<x<1$.

This is easy to check. Solving $f(x,z)=0$ for z, we convert $F(x,z)$ into a function of x alone. Plotting this function for $2\le k\le 30$ and $0<x<1$ shows that the function is always negative, approaching a simple zero at $x=1$; see Fig. 2 for plots of $F(x,z(x))$ as a function of x for some k values. (Checking larger values of k is straightforward, but we stopped at 30.) This eliminates the $n=3$, c_3 case for $k\le 30$.

For $n=2$, set $z=c_2/c_1$. We then have

$$t_1=2c_1c_2-c_1^2=c_1^2(2z-1)$$
$$t_k=c_1c_2^k+c_2c_1^k-c_1^{k+1}=c_1^{k+1}(z^k+z-1),$$

which gives

$$\frac{t_k}{t_1^{(k+1)/2}}=\frac{z^k+z-1}{(2z-1)^{(k+1)/2}}.$$

This function of z is unimodal for $z>1$ (decreasing, then increasing, as can be seen by replacing z with $w=2z-1$) so the maximum of t_k for fixed t_1 occurs at either $c_2=1$ or $c_2=c_1$, resulting in an anticlique or clique. □

We conclude this section with a qualitative feature of phase diagrams for models with constraints on edges and any other simple graph H.

Theorem 4.3. *The phase diagram for an edge-H model is simply connected.*

Proof. This is just the intermediate value theorem. For each fixed edge density ϵ, let g_ϵ be a graphon that maximizes t_H and h_ϵ a graphon which minimizes t_H. Consider the family of graphons $m_{\epsilon,a}(x,y) = ah_\epsilon(x,y) + (1-a)g_\epsilon(x,y)$. When $a = 0$ we get the maximal value of t_H, when $a = 1$ we get the minimal value, and by continuity we have to get all values in between. In other words, the entire phase space between the upper boundary and the lower boundary is filled in. □

Note: This is an application of a technique we learned from Oleg Pikhurko [P].

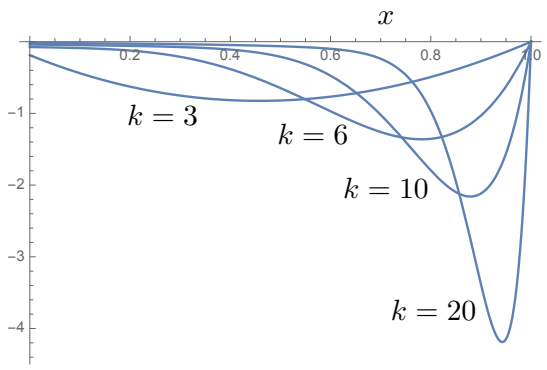

Figure 2: Plots of $F(x, z(x))$ as a function of x for $k = 3$, $k = 6$, $k = 10$ and $k = 20$.

5 Phase Transition for 2-Stars

Theorem 5.1. *For the 2-star model there are inequivalent graphons maximizing the constrained entropy on the line segment* $\{(1/2, \tau_2) \,|\, \tau^* < \tau_2 \leq 2^{-3/2}\}$ *for some* $\tau^* < 2^{-3/2}$. *Moreover, near* $(1/2, 2^{-3/2})$, *the maximizing graphons do not vary continuously with the constraint parameters.*

Proof. For any graphon g, consider the graphon $g'(x,y) = 1 - g(x,y)$. The degree functions for g' and g are related by $d'(x) = 1 - d(x)$. If g has edge and 2-star densities ϵ and τ_2, then g' has edge density $1 - \epsilon$ and 2-star density $\int_0^1 (1 - d(x))^2 dx = 1 - 2\epsilon + \tau_2$. Furthermore, $\mathcal{S}(g') = \mathcal{S}(g)$. This implies that g' maximizes the entropy at $(1-\epsilon, 1-2\epsilon+\tau_2)$ if and only if g maximizes the entropy at (ϵ, τ_2). In particular, if g maximizes the entropy at $(1/2, \tau_2)$, then so does g'. To show that $\mathcal{S}(g)$ has a non-unique maximizer along the upper part of the $\epsilon = 1/2$ line, we must only show that a maximizer g is not related to its mirror g' by a measure-preserving transformation of $[0,1]$.

As noted in Section 4, up to such a transformation there are exactly two graphons corresponding to $(\epsilon, \tau_2) = (1/2, 1/(2\sqrt{2}))$, namely a g-anticlique g_a and a g-clique g_c. These are not related by reordering, since the values of the degree function for the g-clique are $\sqrt{2}/2$ and 0, while those for the g-anticlique are 1 and $1 - \sqrt{2}/2$. Let D be smallest of the following distances in the cut metric: (1) from g_a to g_c, (2) from g_a to the set of symmetric graphons, and (3) from g_c to the set of symmetric graphons.

Lemma 5.2. *There exists $\delta > 0$ such that every graphon with (ϵ, τ_2) within δ of $(1/2, 2^{-3/2})$ is within $D/3$ of either g_a or g_c.*

Proof. Suppose otherwise. Then we could find a sequence of graphons with (ϵ, τ_2) converging to $(1/2, 2^{-3/2})$ that have neither g_a nor g_c as an accumulation point. However, the space of reduced graphons is known to be compact [Lov], so there must be some accumulation point g_∞ that is neither g_a nor g_c. Since convergence in the cut metric implies convergence of the density of all subgraphs, $t_1(g_\infty) = 1/2$ and $t_2(g_\infty) = 2^{-3/2}$. But this contradicts the fact that only g_a and g_c have edge and 2-star densities $(1/2, 2^{-3/2})$. □

By the lemma, no graphon with $\epsilon = 1/2$ and $\tau_2 > 1/2\sqrt{2} - \delta$ is invariant (up to reordering) under $g \to 1 - g$. In particular, the entropy maximizers cannot be symmetric, so there must be two (or more) entropy maximizers, one close to g_a and one close to g_c. □

Moreover, on a path in the parameter space from the anticlique on the upper boundary at $\epsilon = \frac{1}{2} - \delta$ to the clique on the upper boundary at $\epsilon = 1/2 + \delta$, there is a discontinuity in the graphon, where it jumps from being close to g_a to being close to g_c. There must be an odd number of such jumps, and if the path is chosen to be symmetric with respect to the transformation $\epsilon \to 1 - \epsilon$, $\tau_2 \to \tau_2 + 1 - 2\epsilon$, the jump points must be arranged symmetrically on the path. In particular, one of the jumps must be at exactly $\epsilon = 1/2$. This shows that the $\epsilon = 1/2$ line forms the boundary between a region where the optimal graphon is close to g_a and another region where the optimal graphon is close to g_c.

6 Simulations

We now show some numerical simulations in the 2-star model ($\ell = 2, k_1 = 1, k_2 = 2$). Our main aim here is to present numerical evidence that the maximizing graphons in this case are in fact *bipodal*, and to clarify the significance of the degeneracy of Theorem 5.1.

To find maximizing K-podal graphons, we partition the interval $[0, 1]$ into K subintervals $\{I_i\}_{i=1,\dots,k}$ with lengths $c_1, c_2, \cdots, c_K$, that is, $I_i = [c_0 + \cdots + c_{i-1}, c_0 + \cdots + c_i]$ (with $c_0 = 0$). We form a partition of the square $[0, 1]^2$ using the product of this partition with itself. We are interested in functions g that are piecewise constant on the partition:

$$g(x, y) = g_{ij}, \quad (x, y) \in I_i \times I_j, \quad 1 \le i, j \le K, \tag{37}$$

with $g_{ij} = g_{ji}$. We can then verify that the entropy density $\mathcal{S}(g)$, the edge density $t_1(g)$ and the 2-star density $t_2(g)$ become respectively

$$\mathcal{S}(g) = -\frac{1}{2} \sum_{1 \le i,j \le K} [g_{ij} \log g_{ij} + (1 - g_{ij}) \log(1 - g_{ij})] c_i c_j, \tag{38}$$

$$t_1(g) = \sum_{1 \le i,j \le K} g_{ij} c_i c_j, \qquad t_2(g) = \sum_{1 \le i,j,k \le K} g_{ik} g_{kj} c_i c_j. \tag{39}$$

Our objective is to solve the following maximization problem:

$$\max_{\{c_j\}_{1\le j\le K},\{g_{i,j}\}_{1\le i,j\le K}} \mathcal{S}(g), \quad \text{subject to:} \quad t_1(g)=\epsilon, \quad t_2(g)=\tau_2, \quad \sum_{1\le j\le K} c_j = 1, \quad g_{ij}=g_{ji}. \tag{40}$$

We developed in [RRS] computational algorithms for solving this maximization problem and have benchmarked the algorithms with theoretically known results. For a fixed $\tau \equiv (\epsilon, \tau_2)$, our strategy is to first maximize for a fixed number K, and then maximize over the number K. Let $s^K_{(\epsilon,\tau_2)}$ be the maximum achieved by the graphon g_K, then the maximum of the original problem is $s_{(\epsilon,\tau_2)} = \max_K\{s^K_{(\epsilon,\tau_2)}\}$. Our computational resources allow us to go up to $K = 16$ at this time. See [RRS] for more details on the algorithms and their benchmark with existing results.

The most important numerical finding in this work is that, for every pair (ϵ, τ_2) in the interior of the phase space, the graphons that maximize $\mathcal{S}(g)$ are *bipodal*. We need only four parameters (c_1, g_{11}, g_{12} and g_{22}) to describe bipodal graphons (due to the fact that $c_2 = 1-c_1$ and $g_{12} = g_{21}$). For maximizing bipodal graphons, we need only three parameters, since (16) implies that

$$\left(\frac{1}{g_{11}}-1\right)\left(\frac{1}{g_{22}}-1\right) = \left(\frac{1}{g_{12}}-1\right)^2, \tag{41}$$

which was used in our numerical algorithms to simplify the calculations.

We show in Fig. 3 maximizing graphons at some typical points in the phase space. The (ϵ, τ_2) pairs for the plots are respectively: $(0.3, 0.16844286)$ and $(0.3, 0.10339268)$ for the first column (top to bottom), $(0.5, 0.32455844)$ and $(0.5, 0.27485281)$ for the second column, and $(0.7, 0.56270313)$ and $(0.7, 0.50339268)$ for the third column.

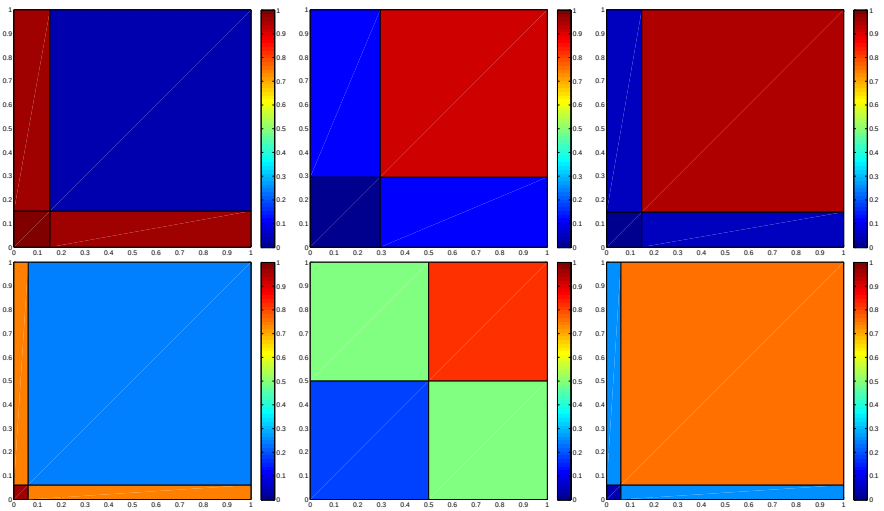

Figure 3: Maximizing graphons at $\epsilon = 0.3$ (left column), $\epsilon = 0.5$ (middle column) and $\epsilon = 0.7$ (right column). For each column τ_2 values decrease from top to bottom.

The values of s corresponding to the maximizing graphons are shown in the left plot of Fig. 4 for a fine grid of (ϵ, σ^2) (with $\sigma^2 = \tau_2 - \epsilon^2$ as defined in Fig. 1) pairs in the phase

space. We first observe that the plot is symmetric with respect to $\epsilon = 1/2$. The symmetry comes from the fact (see the proof of Theorem 5.1) that the map $g \to 1-g$ takes $\epsilon \to 1-\epsilon$, $\tau_2 \to 1-2\epsilon+\tau_2$ and thus $\sigma^2 \to \sigma^2$. To visualize the landscape of s better in the phase space, we also show the cross-sections of $s_{(\epsilon,\tau_2)}(\epsilon,\sigma^2)$ along the lines $\epsilon_k = 0.05k$, $k = 7,\cdots,13$, in the right plots of Fig. 4.

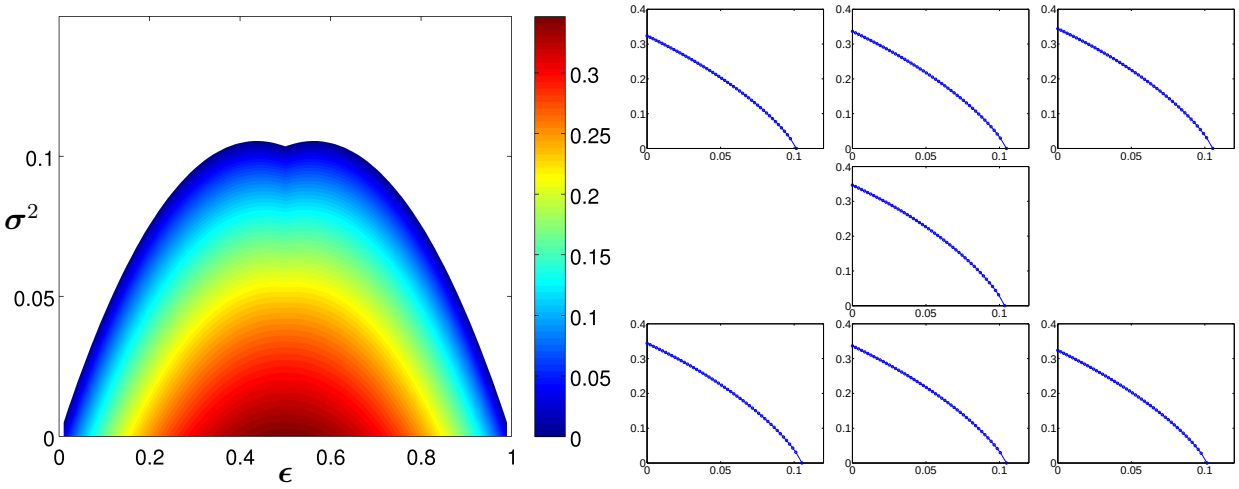

Figure 4: Left: values of $s_{(\epsilon,\tau_2)}$ at different (ϵ,σ^2) pairs; Right: cross-sections of $s_{(\epsilon,\tau_2)}(\epsilon,\sigma^2)$ along lines $\epsilon = \epsilon_k = 0.05k$ $(k = 7,\cdots,13)$ (from top left to bottom right).

We show in the left plot of Fig. 5 the values of c_1 of the maximizing graphons as a function of the pair (ϵ,σ^2). Here we associate c_1 with the set of vertices among V_1 and V_2 that has the larger probability of an interior edge. This is done to avoid the ambiguity caused by the fact that one can relabel V_1, V_2 and exchange c_1 and c_2 to get an equivalent graphon with the same ϵ, τ_2 and $\mathcal{S}$ values. We again observe the symmetry with respect to $\epsilon = 1/2$. The cross-sections of $c_1(\epsilon,\sigma^2)$ along the lines of $\epsilon_k = 0.05k$ $(k = 7,\cdots,13)$ are shown in the right plots of Fig. 5.

The last set of numerical simulations were devoted to the study of a phase transition in the 2-star model. The existence of this phase transition is suggested by the degeneracy in Theorem 5.1. Our numerical simulations indicate that the functions differ to first order in $\epsilon - 1/2$, and that the actual entropy $s_{(\epsilon,\tau_2)} = \max\{s^L{}_{(\epsilon,\tau_2)}, s^R_{(\epsilon,\tau_2)}\}$ has a discontinuity in $\partial_\epsilon s_{(\epsilon,\tau_2)}$ at $\epsilon = 1/2$ above a critical value τ_2^c. Below τ_2^c, there is a single maximizer, of the form

$$g(x,y) = \begin{cases} \frac{1}{2}+\nu & x,y < \frac{1}{2} \\ \frac{1}{2}-\nu & x,y > \frac{1}{2} \\ \frac{1}{2} & \text{otherwise} \end{cases} \tag{42}$$

Here ν is a parameter related to τ_2 by $\tau_2 = 1/4 + \nu^2/4$. Applying the symmetry $g \to 1-g$ and reordering the interval $[0,1]$ by $x \to 1-x$ sends g to itself.

The critical point τ_2^c is located on the boundary of the region in which the maximizer (42) is stable. The value of τ_2^c can be found by computing the second variation of $\mathcal{S}(g)$ within the space of bipodal graphons with fixed values of $(\epsilon = \frac{1}{2}, \tau_2)$, evaluated at the maximizer (42).

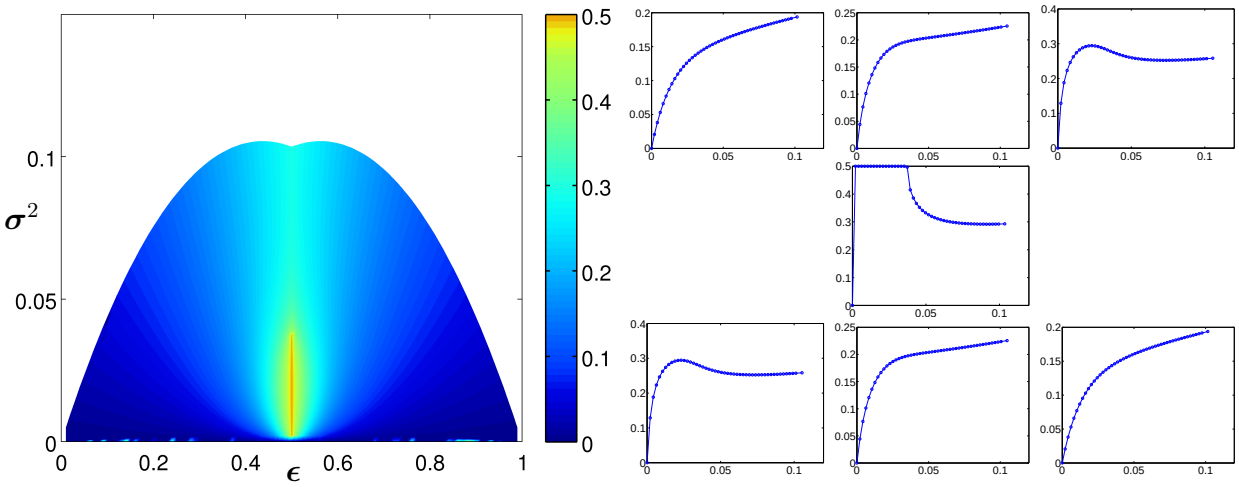

Figure 5: Left: c_1 of the maximizing bipodal graphons as a function of (ϵ, σ^2); Right: cross-sections of $c_1(\epsilon, \sigma^2)$ along lines of $\epsilon_k = 0.05k$ $(k = 7, \cdots, 13)$.

This second variation is positive-definite for ν small (*i.e.* for τ_2 close to $1/4$) and becomes indefinite for larger values of ν. At the critical value of τ_2^c, $\nu = 2\sqrt{\tau_2^c - 1/4}$ satisfies

$$\left(2S(\frac{1}{2}-\nu) - 2S(\frac{1}{2}) + 3\nu S'(\frac{1}{2}-\nu)\right)(2 - \frac{1}{2}S''(\frac{1}{2}-\nu)) + 8\nu^2 S''(\frac{1}{2}-\nu) = 0 \qquad (43)$$

where S' and S'' are respectively the first and second order derivatives of $S(g)$ (defined in (6)) with respect to g. This equation is transcendental, and so cannot be solved in closed form. Solving it numerically for ν leads to the value $\tau_2^c \approx 0.287$, or $\sigma^2 \approx 0.037$. This agrees precisely with what we previously observed in our simulations of optimizing graphons, and corresponds to the point in the left plot Fig. 5 where the $c_1 = 1/2$ region stops.

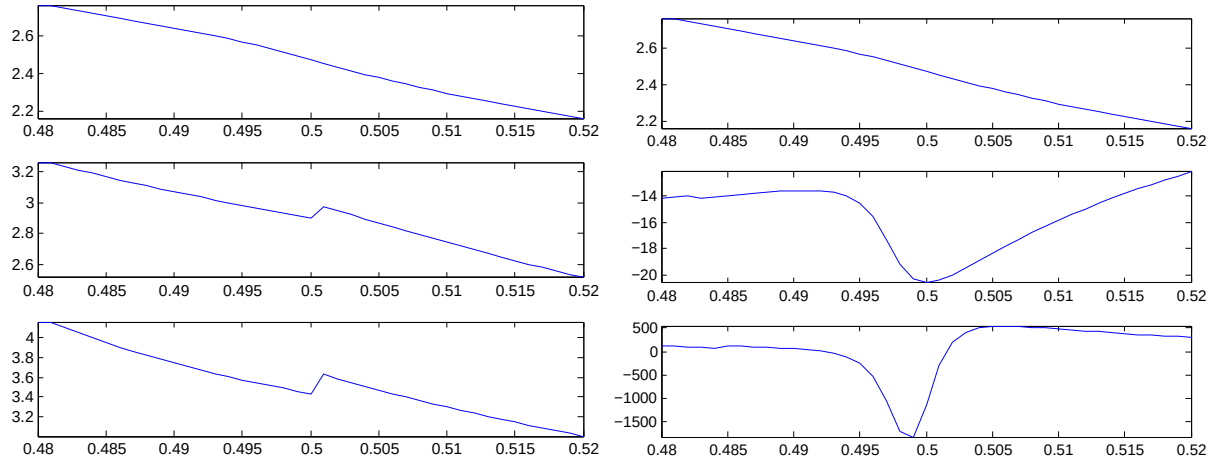

Figure 6: Left: the derivative $\frac{\partial s_{(\epsilon,\tau_2)}}{\partial \epsilon}$ at $\tau_2 = 0.28$ (top), $\tau_2 = 0.30$ (middle) and $\tau_2 = 0.32$ (bottom) in the neighborhood of $\epsilon = 0.5$; Right: the derivatives $\frac{\partial s_{(\epsilon,\tau_2)}}{\partial \epsilon}$ (top), $\frac{\partial^2 s_{(\epsilon,\tau_2)}}{\partial \epsilon^2}$ (middle) and $\frac{\partial^3 s_{(\epsilon,\tau_2)}}{\partial \epsilon^3}$ (bottom) in the neighborhood of $\epsilon = 0.5$ for $\tau_2 = 0.28$.

In the left plot of Fig. 6, we show numerically computed derivatives of $s_{(\epsilon,\tau_2)}$ with respect to ϵ in the neighborhood of $\epsilon = 0.5$ for three different values of τ_2: one below the critical point and two above it. It is clear that discontinuities in the first order derivative of s appears at $\epsilon = 0.5$ for $\tau_2 > \tau_2^c$. When $\tau_2 < \tau_2^c$, we do not observe any discontinuity in the first three derivatives of s.

7 The Taco Model

We next compute the phase space of a model with constrained densities of k-stars for $k = 1, 2, 3$. We call this the "taco" model from the shape of its phase space; see Figure 8. Let t_k be the density of k-stars.

A simple application of the Cauchy-Schwarz inequality applied to $d^{1/2}$ and $d^{3/2}$ shows that

$$t_1 t_3 \geq t_2^2. \tag{44}$$

Under the symmetry $g(x, y) \to 1 - g(x, y)$ we have

$$\begin{aligned} t_1 &\to 1 - t_1, \\ t_2 &\to 1 - 2t_1 + t_2, \\ t_3 &\to 1 - 3t_1 + 3t_2 - t_3. \end{aligned}$$

The inequality (44) thus gives another inequality

$$(1 - t_1)(1 - 3t_1 + 3t_2 - t_3) \geq (1 - 2t_1 + t_2)^2. \tag{45}$$

These two inequalities (or more precisely, the corresponding equalities) form two boundaries of the phase space. The rest of the boundary is determined by an Euler-Lagrange equation, and is given by the locus of values obtained by the tripodal $\{0, 1\}$-valued graphons of Figure 7, right, when x_1, x_2, x_3 vary over all possible values with $x_1 + x_2 + x_3 = 1$.

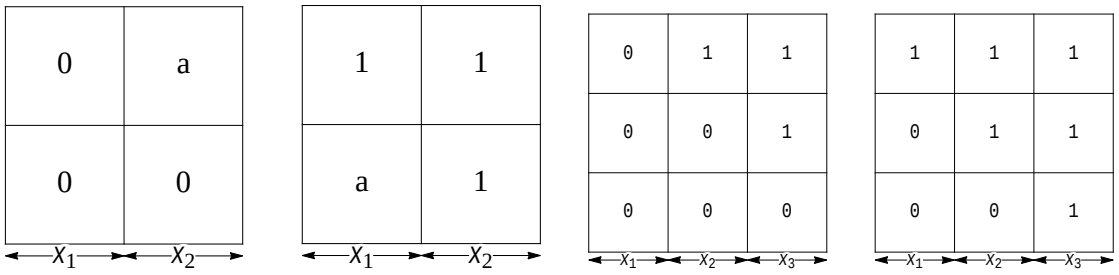

Figure 7: Graphons defining the boundary of the phase space.

Theorem 7.1. *The phase space of t_1, t_2, t_3-values of graphons is the solid domain delimited by these four algebraic surfaces. See Figure 8. Its boundary is piecewise analytic with four pieces; the four families of graphons parameterized by the four pieces are indicated in Figure 7.*

The two upper surfaces (corresponding to the two right graphons in Figure 7) meet along three algebraic curves: one corresponding to $a = 1$ in the first graphon (the curve along the top of the front part of the "shell" of the taco), one corresponding to $a = 0$ in the second graphon (the curve along the top of the back part of the shell) and a third curve running along the center of the taco. Along this third curve we have two inequivalent graphons. This coexistence curve is a high-degree algebraic curve in x_1, x_2, x_3.

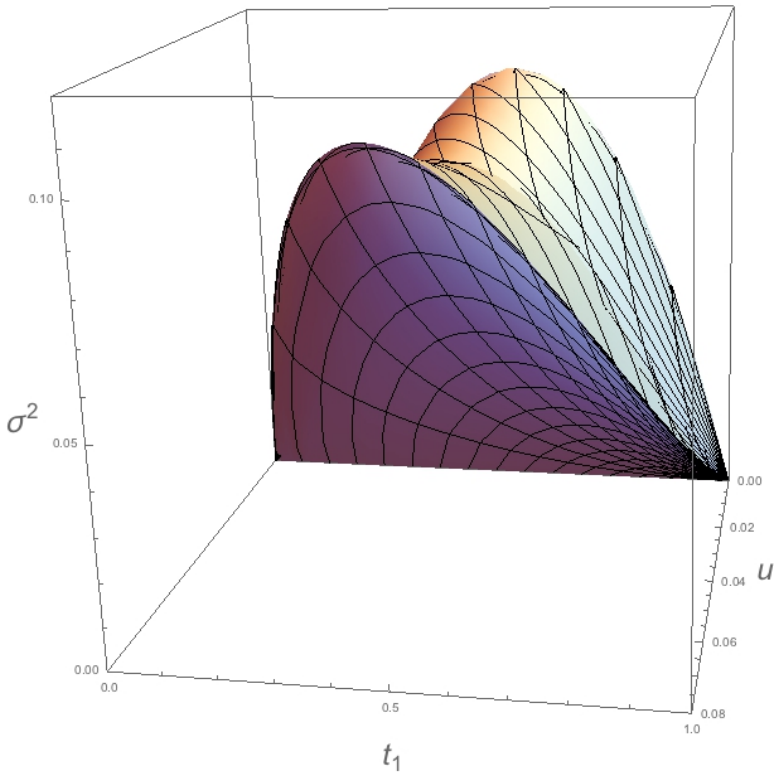

Figure 8: The phase space, or "taco". The axes are $t_1, \sigma^2 = t_2 - t_1^2$, and $u = t_3 - t_2^2/t_1$. The front (dark shaded) face is the surface defined by the equality in (45). The back face (flush with the back boundary $u = 0$) is defined by equality in (44). The "upper" faces are surfaces defined by the tripodal $0-1$-valued graphons of Figure 7, right. Note that if we project the taco along the u direction we get Figure 1, right; one sees that the coexistence line there in fact lifts to two different curves in the taco phase space.

Proof. The two graphon families on the left in Figure 7 correspond to t_1, t_2, t_3-values satisfying equality in (44) and (45), respectively, so they indeed comprise part of the boundary of the phase space. We need to verify only that there are no graphons above the "upper" boundaries.

By Theorem 3.1 we may assume that $d(x)$ is monotone increasing and $g(x, y)$ is monotone increasing in both variables.

We first show that g only takes values in $\{0, 1\}$. To accomplish this, we show first that if $g(x, y) \in (0, 1)$ on a positive measure set, then d can take at most 2 values, then show

that d can in fact take at most one value (which would mean the graphon is an Erdős-Renyi graphon).

Suppose that $g(x,y) \in (0,1)$ on an open set containing a rectangle R_1, and that d takes at least three values. We claim that $d(x)$ and $d(y)$ are constant on $R_1 = [x', x''] \times [y', y'']$. If $d(x)$ takes at least three values globally, and two values on R_1, we can find $x_1, x_2 \in [x', x'']$ and $x_3 \in [0,1]$ such that $d(x_1) < d(x_2) < d(x_3)$ or $d(x_3) < d(x_1) < d(x_2)$. Suppose we are in the first case; the second is similar. By shifting mass δ from $g(x_1, y)$ to $g(x_2, y)$ and ε from $g(x_3, y)$ to $g(x_2, y)$, where

$$\varepsilon = \frac{d(x_2)^2 - d(x_1)^2}{d(x_3)^2 - d(x_2)^2}\delta + O(\delta)^2$$

we can fixing t_1, t_3 while increasing t_2; a short calculation yields

$$\Delta t_2 = \frac{(d(x_2) - d(x_1))(d(x_3) - d(x_1))}{d(x_2) + d(x_3)}\delta + O(\delta^2).$$

This contradicts maximality.

Suppose that $g(x,y) \in (0,1)$ on R_1 and $d(x)$ is constant on R_1. Let $(x_1, y_1), (x_2, y_1) \in R_1$ and find a point $(x_3, y_1) \notin R_1$ with $d(x_3) \neq d(x_1)$. We assume $d(x_3) > d(x_1)$, the other case is similar. Since $x_3 > x_1$ by monotonicity, $g(x_3, y_1) > 0$. We move mass ε from (x_3, y_1) to (x_1, y_1), and mass δ from (x_2, y_1) to (x_1, y_1). Here we let ε *be of order* δ^2. The changes in t_3 and t_2 are then respectively

$$\begin{aligned} \Delta t_3 &= 3\varepsilon(-d(x_3)^2 + d(x_1)^2) + 6\delta^2 d(x_1) + O(\delta^3) \\ \Delta t_2 &= 2\varepsilon(-d(x_3) + d(x_1)) + 2\delta^2. \end{aligned}$$

In particular by setting

$$\varepsilon = \frac{2d(x_1)\delta^2}{d(x_3)^2 - d(x_1)^2} + O(\delta^3)$$

we can fix t_3 while increasing t_2: with this value of ε,

$$\Delta t_2 = \frac{2(d(x_3) - d(x_1))\delta^2}{d(x_3) + d(x_1)}.$$

This contradicts maximality of t_2.

This argument show in fact that either d takes only *one* value, or g takes values only in $\{0,1\}$. Thus we may assume g takes values only in $\{0,1\}$.

We now assume that g takes values only in $\{0,1\}$, and, reversing our convention for the rest of the proof, is monotone *decreasing* in both variables.

As in the proof of Theorem 4.1, we approximate the boundary curve between $g = 0$ and $g = 1$ with a zig-zag path. We can reduce the number of zigs and zags as before, except now we must keep t_1 and t_3 fixed, and increase t_2 along the way.

Take four consecutive points on the path: $(c_{i-1}, c_j)(c_i, c_j)(c_i, c_{j-1})(c_{i+1}, c_{j-1})$, lying above or on the diagonal $x = y$ and none of which is $(0,1)$. We will vary the three values c_i, c_j, c_{j-1}. (The case of a vertical, horizontal, vertical subpath is similar).

As we saw in the proof of Theorem 4.1, by varying c_i, c_j, c_{j-1} we can always increase t_2; in fact there is a two-parameter family of ways to increase t_2: varying just c_i, c_j or just c_i, c_{j-1}, for example.

By Theorem 3.1, t_2 is locally critical (under changes in c_i, c_j, c_{j-1}) if and only if the quantities $c_i - c_{i-1}, c_j - c_{j-1}$, and $c_{i+1} - c_i$ are equal. Similarly one finds that t_3 is locally critical if and only if

$$(c_i - c_{i-1})(2c_i + c_{i-1}) = (c_j - c_{j-1})(2c_j + c_{j-1}) \tag{46}$$

and

$$(c_{i+1} - c_i)(2c_i + c_{i+1}) = (c_j - c_{j-1})(2c_{j-1} + c_j). \tag{47}$$

If t_3 is not critical there is a two-parameter family of perturbations of t_3, and thus a one-parameter family fixing t_3. Hence, by varying in one direction or another along this family it is possible to increase t_2 while holding t_3 fixed. We thus need only consider the case when t_3 is critical. In this case any perturbation will increase t_3 to second order, but t_2 to first order, since (46) and (47) imply that none of $c_i - c_{i-1}, c_j - c_{j-1}$, and $c_{i+1} - c_i$ are equal.

This is sufficient for our purposes since the boundary (45) is not vertical in the direction of t_3. We conclude that t_2 is maximized only when there are no four consecutive points as above, that is, g (reordered to be increasing) has one of the two forms of Figure 7, right.

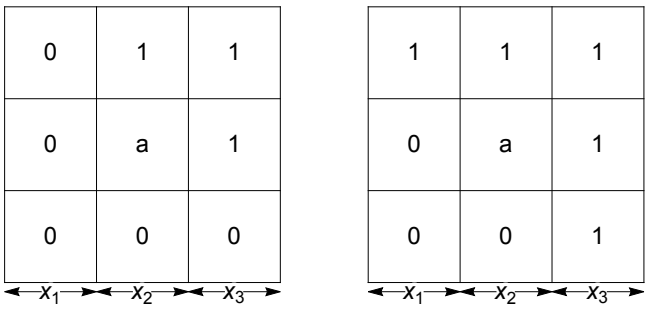

Figure 9: Graphons covering the interior of the taco.

It remains to show that the taco is full, that is, each interior point is realizable. We consider the 3-parameter family Ω of tripodal graphons of Figure 9 (left). This family maps into t_1, t_2, t_3-space with Jacobian nonzero on the interior: the Jacobian is

$$(a-1)^2 x_2^4 \left(a^2 x_2^3 + 2a^2 x_3 x_2^2 + 2a x_3^2 x_2 + x_3 x_2^2 + x_3^2 x_2 + x_3^3\right)$$

which is a sum of positive terms. Moreover the boundary of Ω maps to the part of the taco boundary consisting of three of the faces of the taco, only missing an area below one of the upper faces (see Figure 10). The map is thus a homeomorphism on the interior of Ω to the interior of its image. Now let Ω' be the image of Ω under the involution, or equivalently, the graphons from Figure 9 (right); this covers the remaining part of the taco. □

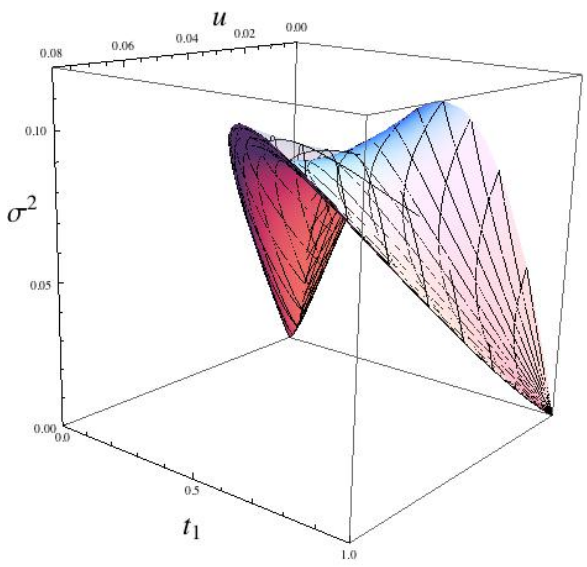

Figure 10: The image of the upper boundary of Ω; shown in transparent rendering is the other half of the upper boundary of the taco, which is part of the boundary of the image of Ω'.

8 A Finitely Forced Model

We have shown that, in the interior of the phase space, entropy maximizers with edge and k-star densities as constraints are multipodal. It is known from extremal graph theory that this is not true in general on the boundary of the phase space. We now briefly look at this issue using the concept of finitely forcible graphons introduced in [LS3].

Let $h(x, y)$ be any doubly monotonic function with 0 as a regular value, and consider the graphon

$$g(x,y) = \begin{cases} 1 & h(x,y) > 0 \\ 0 & h(x,y) < 0 \end{cases}. \tag{48}$$

Then it is shown in [LS3] that for this graphon, the density of the signed quadrilateral subgraph Q (with signs going around the quadrilateral as $+$, $-$, $+$, $-$), is zero. In other words,

$$t_Q(g) = \int g(w,x)[1-g(x,y)]g(y,z)[1-g(w,z)]dw\,dx\,dy\,dz = 0, \tag{49}$$

where we have labelled the four vertices as w, x, y, z. It is straightforward to verify that $t_Q = t_1^2 - 2\tilde{t}_3 + \tilde{t}_4$ where $\tilde{t}_3$ is the density of 3-chains while $\tilde{t}_4$ is the density of the un-signed quadrilateral:

$$\begin{array}{rcl} \tilde{t}_3(g) & = & \int g(w,x)g(x,y)g(y,z)dw\,dx\,dy\,dz, \\ \tilde{t}_4(g) & = & \int g(w,x)g(x,y)g(y,z)g(z,w)dw\,dx\,dy\,dz. \end{array} \tag{50}$$

The triangle graphon g_T defined by

$$g_T(x,y) = \begin{cases} 1 & x+y>1 \\ 0 & x+y<1 \end{cases} \tag{51}$$

is a special case of (48) with $h(x,y) = x+y-1$. For the triangle graphon we check that it has edge density $t_1 = 1/2$, 2-star density $t_2 = 1/3$, 3-chain density $\tilde{t}_3 = 5/24$ and quadrilateral density $\tilde{t}_4 = 1/6$. This clearly gives us $t_Q(g_T) = 0$.

It is shown in [LS3] that (up to rearranging vertices) graphons of the form (48) are the only kind of graphon for which $t_Q(g) = 0$. Moreover, among graphons of the form (48) the density of edges minus the density of 2-stars is at most 1/6, and this upper bound is achieved uniquely by the graphon g_T. Thus the triangle graphon g_T is finitely forcible with the two constraints:

1. $\zeta_1(g) \equiv t_1(g) - t_2(g) - \frac{1}{6} \geq 0$, i.e. the density of edges should be 1/6 greater than the density of 2-stars.
2. $\zeta_2(g) \equiv t_Q(g) = t_1^2(g) - 2\tilde{t}_3(g) + \tilde{t}_4(g) = 0$, i.e. the density of the signed quadrilateral subgraph Q should be zero.

Here we look at a path toward the triangle graphon by considering the parameterized family of graphons $g_\alpha(x,y) = \alpha + (1-2\alpha)g_T(x,y)$ $(0 \leq \alpha \leq 0.5)$. We attempt to maximize the entropy among graphons that have the same values of $t = (t_1, t_2, \tilde{t}_3, \tilde{t}_4)$ as g_α. We first check that

$$\begin{aligned} t_1 &= 1/2 \\ t_2 &= (1-\alpha+\alpha^2)/3 \\ \tilde{t}_3 &= (5-8\alpha+8\alpha^2)/24 \\ \tilde{t}_4 &= (1-3\alpha+5\alpha^2-4\alpha^3+2\alpha^4)/6. \end{aligned}$$

This gives $\zeta_1(g_\alpha) = \alpha(1-\alpha)/3$ and $\zeta_2(g_\alpha) = \alpha(1+\alpha-4\alpha^2+2\alpha^3)/6$.

We can show that by enforcing the densities $(t_1, t_2, \tilde{t}_3, \tilde{t}_4)$, we are approaching the triangle graphon from the interior of the profile when we let $\alpha \to 0$, as stated in the following theorem.

Theorem 8.1. *For any $\alpha > 0$, the values of t lie in the interior of the profile.*

Proof. Since the graphon g_a is strictly between 0 and 1, it is enough to show that the four functional derivatives, $\delta t_1/\delta g$, $\delta t_2/\delta g$, $\delta \tilde{t}_3/\delta g$ and $\delta \tilde{t}_4/\delta g$, are linearly independent functions of x and y, since then by varying g we can change t in any direction to first order. A simple computation shows that $\delta t_1/\delta g(x,y)$ is a constant, $\delta t_2/\delta g(x,y)$ is a linear polynomial in x and y, and $\delta \tilde{t}_3/\delta g(x,y)$ is a quadratic polynomial with an xy term as well as linear terms. These three are analytic and manifestly linearly independent.

However, $\delta \tilde{t}_4/\delta g(x,y)$ is not analytic across the line $x+y=1$, and so cannot be a linear combination of the first three functional derivatives. To see this, it is enough to consider the case of $\alpha = 0$. $\delta t_4/\delta g(x,y)$ is a multiple of the probability of x being connected to y via

a 3-chain y-z-w-x. When $x+y<1$, this is exactly xy, since if $z>1-y$ and $w>1-x$ then $z+w$ is automatically greater than 1. When $x+y>1$, the requirement that $z+w>1$ provides an extra condition, and the functional derivative is strictly less than xy. □

We can now show that if we try to fit the density t with M-podal graphons, then M blows up as α goes to zero. More precisely,

Theorem 8.2. *For each positive integer M there is an $\epsilon_M>0$ such that for $\alpha<\epsilon_M$ there are no M-podal graphons whose densities t are the same as those of g_α.*

Proof. Suppose otherwise. Then we could find M-podal graphons for arbitrarily small α. Since the space of M-podal (or smaller) graphons is compact, we can find a subsequence that converges to g_0 as $\alpha\to 0$. But densities vary continuously with the graphon, being simply integrals of products of g's. So the densities of the M-podal graphon g_0 are the same as the densities of g_T. But that is a contradiction, since g_T was finitely forced. □

Therefore, we know that $M\to\infty$ as $\alpha\to 0$. The fact that the densities are polynomials in g (of degree at most 4) suggests that the growth should be at least a power law.

We performed two sets of numerical simulations. In the first set of simulations, we enforce the densities $(t_1,t_2,\tilde{t}_3,\tilde{t}_4)$ by solving the following minimization problem for some α:

$$\max_{\{c_j\}_{1\le j\le K},\{g_{i,j}\}_{1\le i,j\le K}} \mathcal{S}(g),\quad \text{subject to:}\quad t(g)=(\tau_1,\tau_2,\tilde{\tau}_3,\tilde{\tau}_4),\quad \sum_{1\le j\le K} c_j=1,\quad g_{ij}=g_{ji}. \tag{52}$$

For values of $\alpha\in[0.001,0.5)$ we get multipodal maximizers with a small number of podes. To be precise we obtain, numerically, 3-podal maximizers for α values in $(0.02,0.5)$, 4-podal maximizers for α values in $(0.004,0.020)$, and 5-podal maximizers for α values in $(0.001,0.004)$. The transition from 3-podal to 4-podal occurs around $\alpha=0.02$, and the transition from 4-podal to 5-podal occurs around $\alpha=0.004$; see the top row of Fig. 11 for typical 3-, 4- and 5-podal maximizers we obtained in this case.

In the second set of simulations, we solve a similar minimization problem that enforce the constaints on ζ_1 and ζ_2, instead of the four densities. For α values in $[0.001,0.5)$ we again get multipodal maximizers with a small number of podes. Precisely, we obtain 2-podal maximizers for α values in $(0.04,0.5)$, 3-podal maximizers for α values in $(0.015,0.040)$, and 4-podal maximizers for α values in $(0.001,0.015)$. The transition from 2-podal to 3-podal occurs around $\alpha=0.04$, and the transition from 3-podal to 4-podal occurs around $\alpha=0.015$; see the bottom row of Fig. 11 for typical 2-, 3- and 4-podal maximizers we obtained in this case.

Overall, our numerical simulations show that in a large fraction of the profile the maximizing graphons are multipodal with a small number of podes. The simulations also demonstrate that M increases as α decreases. However, the numerical evidences are far from conclusive in the sense that we are not able to push α small enough to see the (necessary) blow up behavior of M more precisely, let alone the nature of the optimizing graphon as that occurs.

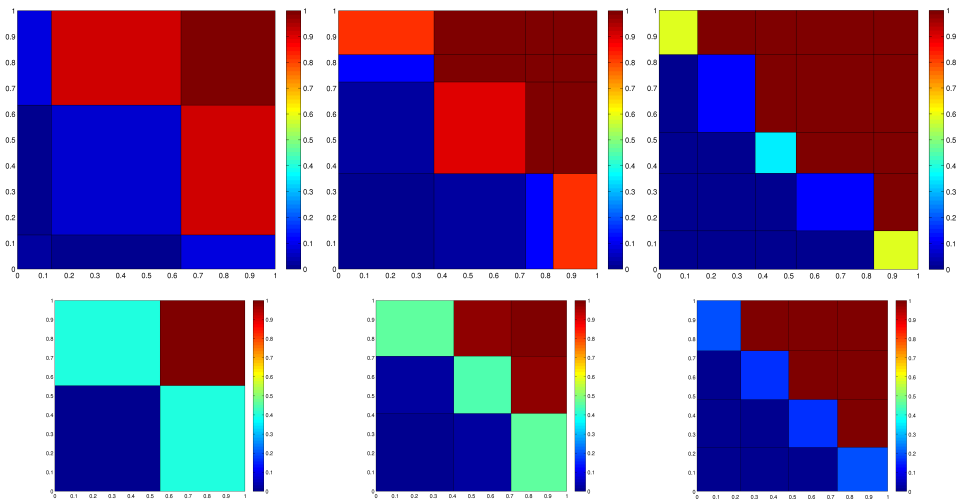

Figure 11: Maximizing graphons for the finitely forced model. Top row: maximizers with $(t_1, t_2, \tilde{t}_3, \tilde{t}_4)$ constraints at $\alpha = 0.0200$ (left, 3-podal), $\alpha = 0.0050$ (middle, 4-podal) and $\alpha = 0.0015$ (right, 5-podal); Bottom row: maximizers with (ζ_1, ζ_2) constraints at $\alpha = 0.0200$ (left, 2-podal), $\alpha = 0.0050$ (middle, 3-podal) and $\alpha = 0.0015$ (right, 4-podal).

9 Conclusion

We first compare our results with exponential random graph models (ERGMs), also based on given subgraph densities; see [CD, RY, LZ, Y, YRF, AZ] for previous mathematical work on their asymptotics. For this we contrast the basic optimization problems underlying ERGMs and the models of this paper.

Intuitively the randomness in such random graph models arises, in modeling large networks, by starting with an assumption that a certain set of subgraphs $H = (H_1, \ldots, H_\ell)$ are 'significant' for the networks; one can then try to understand a large network as a 'typical' one for certain values $t_H(g) = (t_{H_1}, \ldots, t_{H_\ell})$ of the densities t_{H_j} of those subgraphs. Large deviations theory [CV] can then give probabilistic descriptions of such typical graphs through a variational principle for the constrained Shannon entropy, $s_\tau = \sup_{g|t_H(g)=\tau} \mathcal{S}(g)$.

In this paper, as in [RS1, RS2, RRS], we use such constrained optimization of entropy, and by analogy with statistical mechanics we call such models 'microcanonical'. In contrast, ERGMs are analogues of 'grand canonical' models of statistical mechanics. As noted in Section 2, our microcanonical version consists of maximizing $\mathcal{S}(g)$ over graphons g with fixed values $t_H(g) = \tau$, leading to a constrained-maximum entropy $s_\tau = \sup_{g|t_H(g)=\tau} \mathcal{S}(g)$. The optimizing graphons satisfy the Euler-Lagrange variational equation, $\delta[\mathcal{S}(g) + \beta \cdot t_H(g)] = 0$, together with the constraints $t_H(g) = \tau$, for some set of Lagrange multipliers $\beta = (\beta_1, \ldots, \beta_\ell)$.

For the ERGM (grand canonical) approach, instead of fixing $t_H(g)$ one *maximizes* $\tilde{F}(g) = \mathcal{S}(g) + \beta \cdot t_H(g)$ for fixed β, obtaining

$$F_\beta = \sup_g \tilde{F}(g) = \sup_g [\mathcal{S}(g) + \beta \cdot t_H(g)]. \tag{53}$$

It is typical for there to be a loss of information in the grand canonical modelling of large graphs. One way to see the loss is by comparing the parameter ("phase") space $\Sigma_{mc} = \{\tau\}$ of the microcanonical model with that for the grand canonical model, $\Sigma_{gc} = \{\beta\}$. For each point β of Σ_{gc} there are optimizing graphons $\tilde{g}_\beta$ such that $\tilde{F}(\tilde{g}_\beta) = F_\beta$, and for each point τ of Σ_{mc} there are optimizing graphons $\tilde{g}_\tau$ such that $t_H(\tilde{g}_\tau) = \tau$ and $\mathcal{S}(\tilde{g}_\tau) = s_\tau$. Defining τ' as $t_H(\tilde{g}_\beta)$ it follows that $\tilde{g}_\beta$ maximizes $\mathcal{S}(g)$ under some constraint τ, namely $\mathcal{S}(\tilde{g}_\beta) = s_{\tau'}$. But the converse fails: there are some τ for which no optimizing $\tilde{g}_\beta$ satisfies $t_H(\tilde{g}_\beta) = \tau$ [CD, RS1].

This asymmetry is particularly acute for the k-star models we discuss in this paper: it follows from [CD] that all of Σ_{gc} is represented only on the lower boundary curve of Σ_{mc}, $\tau_k = \epsilon^k$: see Fig. 1. If one is interested in the influence of certain subgraph densities in a large network it is therefore preferable to use constrained optimization of entropy rather than to use the ERGM approach.

Finally, in trying to understand the 'phase transition' in the ERGM edge/triangle model it seems significant that the functional derivatives of the densities $\delta t/\delta g$ are linearly dependent at the optimizing (constant) graphons relevant to that transition. This was quite relevant in the perturbative analysis along the Erdős-Rényi curve in the microcanonical edge/triangle model [RS1]. And of course when the $\delta t/\delta g$ are linearly dependent they cannot play their usual role as coefficients in the expansion of the entropy s_τ.

Next we consider the role of multipodal states in modeling large graphs. In [RS1, RS2, RRS] evidence, but not proof, of multipodal entropy optimizers was found throughout the phase space of the microcanonical edge/triangle model, and in this paper we have proven this to hold throughout the phase space of all k-star models. Consider more general microcanonical graph models with constraints on edge density, $e(g)$, and the densities $t_H(g)$ of a finite number of other subgraphs, H. We are interested in the generality of multipodality for entropy maximizing graphons in such models. As noted in the Introduction there are known examples (in some sense 'rare': see Theorem 7.12 in [LS3]) with nonmultipodality on phase space boundaries, but this is not known to occur in the interior of any phase space.

To pursue this we first note a superficial similarity between the subject of extremal graphs and the older subject of 'densest packings of bodies': Given a finite collection $\mathcal{C}$ of bodies $\{B_1, \ldots, B_\ell\}$ in $\mathcal{R}^d$ determine those nonoverlapping arrangements, of unlimited numbers of congruent copies of the B's, which maximize the fraction of $\mathcal{R}^d$ covered by the B's. (See [Fej] for an overview.) As in extremal graph theory few examples have been solved, the main ones being congruent spheres for dimensions $d \leq 3$ and those bodies which can tile space, such as congruent regular hexagons in the plane. Based on this limited experience the assumption/expectation developed that for every collection $\mathcal{C}$ there would be a 'crystalline' densest packing, a packing whose symmetry group was small (cocompact) in the group of symmetries of $\mathcal{R}^d$. This assumption was proven incorrect in 1966 by the construction of 'aperiodic tilings'; see [Sen, Ra] for an overview. One can therefore draw a parallel between aperiodic 'counterexamples' in the study of densest packings, and nonmultipodal 'counterexamples' in extremal graph theory. Using nonoverlapping bodies to model molecules, physicists have applied the formalism of statistical mechanics to packings of bodies. Packings of spheres then give rise to the 'hard sphere model' which is a simple model for which simulation (not

proof) shows the emergence of a crystalline phase in the interior of the phase space [Low]. More recently, aperiodic tilings have been used to model quasicrystalline phases of matter. (See [J] for an introductory guide to quasicrystals.) Although it is expected that the tilings, corresponding to optimal density on the boundary of the microcanonical phase space, give rise to an emergent quasicrystalline phase in the interior, there is much less simulation evidence of this, as yet, than for crystalline phases emerging from crystalline sphere packings; see [AR] and references therein.

Getting back to networks we note that simple constraints give rise to multipodal optimal graphs on the boundary of the phase space, and also [RS1, RS2, RRS] multipodal phases in the interior, in parallel to the crystalline situation in packing. By analogy with packing therefore, a natural question is: do the nonmultipodal 'counterexamples' on the phase space boundary of random graph models give rise to nonmultipodal *phases* in the interior of the phase space, in parallel to aperiodic tilings and quasicrystalline phases? (As in statistical mechanics a phase is defined as a connected open subset of the phase space of the model, in which the entropy is analytic; see [RS1].)

Question 1. Are random graph phases always multipodal?

Our attempt to investigate this in Section 8 was inconclusive.

Multipodality is a useful tool in understanding phases. For instance in the edge/triangle model [RS1, RS2, RRS] even a cursory inspection of the largest values of such an optimizing graphon concentrates attention on the conditions under which edges tend to clump together (fluid-like behavior) or push apart into segregated patterns (solid-like behavior). More specifically, we note that in simulations of the edge/triangle model [RRS] it is very noticeable that at densities above the Erdős-Rényi curve the optimizing graphons are always monotone, while this is rarely if ever the case below the curve. We proved in this paper that in k-star models, for which densities are *always* above the ER curve, the optimizing graphons are always monotone. Consider a general model with two densities, edges and some graph H. The ER curve is $\tau_H = \epsilon^k$ where k is the number of edges in H. A natural question is:

Question 2. Are the optimizing graphons always monotone above the ER curve in such random graph models?

In equilibrium statistical mechanics [Ru] one can rarely understand directly the equilibrium distribution in a useful way, at least away from extreme values of energy or pressure, so one determines the basic characteristics of a model by estimating order parameters or other secondary quantities. In random graph models multipodal structure of the optimizing state gives hope for a more direct understanding of the emergent properties of a model. This would be a significant shift of viewpoint.

Acknowledgments

The authors gratefully acknowledge useful discussions with Mei Yin and references from Miki Simonovits, Oleg Pikhurko and Daniel Kral. The computational codes involved in this

research were developed and debugged on the computational cluster of the Mathematics Department of UT Austin. The main computational results were obtained on the computational facilities in the Texas Super Computing Center (TACC). We gratefully acknowledge this computational support. R. Kenyon was partially supported by the Simons Foundation. This work was also partially supported by NSF grants DMS-1208191, DMS-1208941, DMS-1321018 and DMS-1101326.

References

[AK] R. Ahlswede and G.O.H. Katona, Graphs with maximal number of adjacent pairs of edges, Acta Math. Acad. Sci. Hungar. 32 (1978) 97-120

[AR] D. Aristoff and C. Radin, First order phase transition in a model of quasicrystals, J. Phys. A: Math. Theor. 44(2011), 255001.

[AZ] D. Aristoff and L. Zhu, On the phase transition curve in a directed exponential random graph model, arxiv:1404.6514

[B] B. Bollobas, Extremal graph theory, Dover Publications, New York, 2004.

[BCL] C. Borgs, J. Chayes and L. Lovász, Moments of two-variable functions and the uniqueness of graph limits, Geom. Funct. Anal. 19 (2010) 1597-1619.

[BCLSV] C. Borgs, J. Chayes, L. Lovász, V.T. Sós and K. Vesztergombi, Convergent graph sequences I: subgraph frequencies, metric properties, and testing, Adv. Math. 219 (2008) 1801-1851.

[CD] S. Chatterjee and P. Diaconis, Estimating and understanding exponential random graph models, Ann. Statist. 41 (2013) 2428-2461.

[CDS] S. Chatterjee, P. Diaconis and A. Sly, Random graphs with a given degree sequence, Ann. Appl. Probab. 21 (2011) 1400-1435.

[CV] S. Chatterjee and S.R.S. Varadhan, The large deviation principle for the Erdős-Rényi random graph, Eur. J. Comb. 32 (2011) 1000-1017.

[Fej] L. Fejes Tóth, Regular Figures, Macmillan, New York, 1964.

[J] C. Janot, Quasicrystals: A primer, Oxford University Press, Oxford, 1997.

[Lov] L. Lovász, Large networks and graph limits, American Mathematical Society, Providence, 2012.

[Low] H. Löwen, Fun with hard spheres, In: "Spatial Statistics and Statistical Physics", edited by K. Mecke and D. Stoyan, Springer Lecture Notes in Physics, volume 554, pages 295–331, Berlin, 2000.

[LS1] L. Lovász and B. Szegedy, Limits of dense graph sequences, J. Combin. Theory Ser. B 98 (2006) 933-957.

[LS2] L. Lovász and B. Szegedy, Szemerédi's lemma for the analyst, GAFA 17 (2007) 252-270.

[LS3] L. Lovász and B. Szegedy, Finitely forcible graphons, J. Combin. Theory Ser. B 101 (2011) 269-301.

[LZ] E. Lubetzky and Y. Zhao, On replica symmetry of large deviations in random graphs, Random Structures and Algorithms (to appear), arXiv:1210.7013 (2012).

[N] M.E.J. Newman, Networks: an Introduction, Oxford University Press, 2010.

[P] O. Pikhurko, private communication.

[R] C. Reiher, The clique density theorem, arXiv:1212.2454.

[Ra] C. Radin, Miles of Tiles, Student Mathematical Library, Vol 1, Amer. Math. Soc., Providence, 1999.

[RRS] C. Radin, K. Ren and L. Sadun, The asymptotics of large constrained graphs, J. Phys. A: Math. Theor. 47 (2014) 175001.

[RS1] C. Radin and L. Sadun, Phase transitions in a complex network, J. Phys. A: Math. Theor. 46 (2013) 305002.

[RS2] C. Radin and L. Sadun, Singularities in the entropy of asymptotically large simple graphs, arXiv:1302.3531 (2013).

[Ru] D. Ruelle Statistical Mechanics; Rigorous Results, Benjamin, New York, 1969.

[RY] C. Radin and M. Yin, Phase transitions in exponential random graphs, Ann. Appl. Probab. 23 (2013) 2458-2471.

[Sen] M. Senechal, Quasicrystals and geometry, Cambridge University Press, Cambridge, 1995.

[TET] H. Touchette, R.S. Ellis and B. Turkington, Physica A 340 (2004) 138-146.

[Y] M. Yin, Critical phenomena in exponential random graphs, J. Stat. Phys. 153 (2013) 1008-1021.

[YRF] M. Yin, A. Rinaldo and S. Fadnavis, Asymptotic quantization of exponential random graphs, arXiv:1311.1738 (2013).